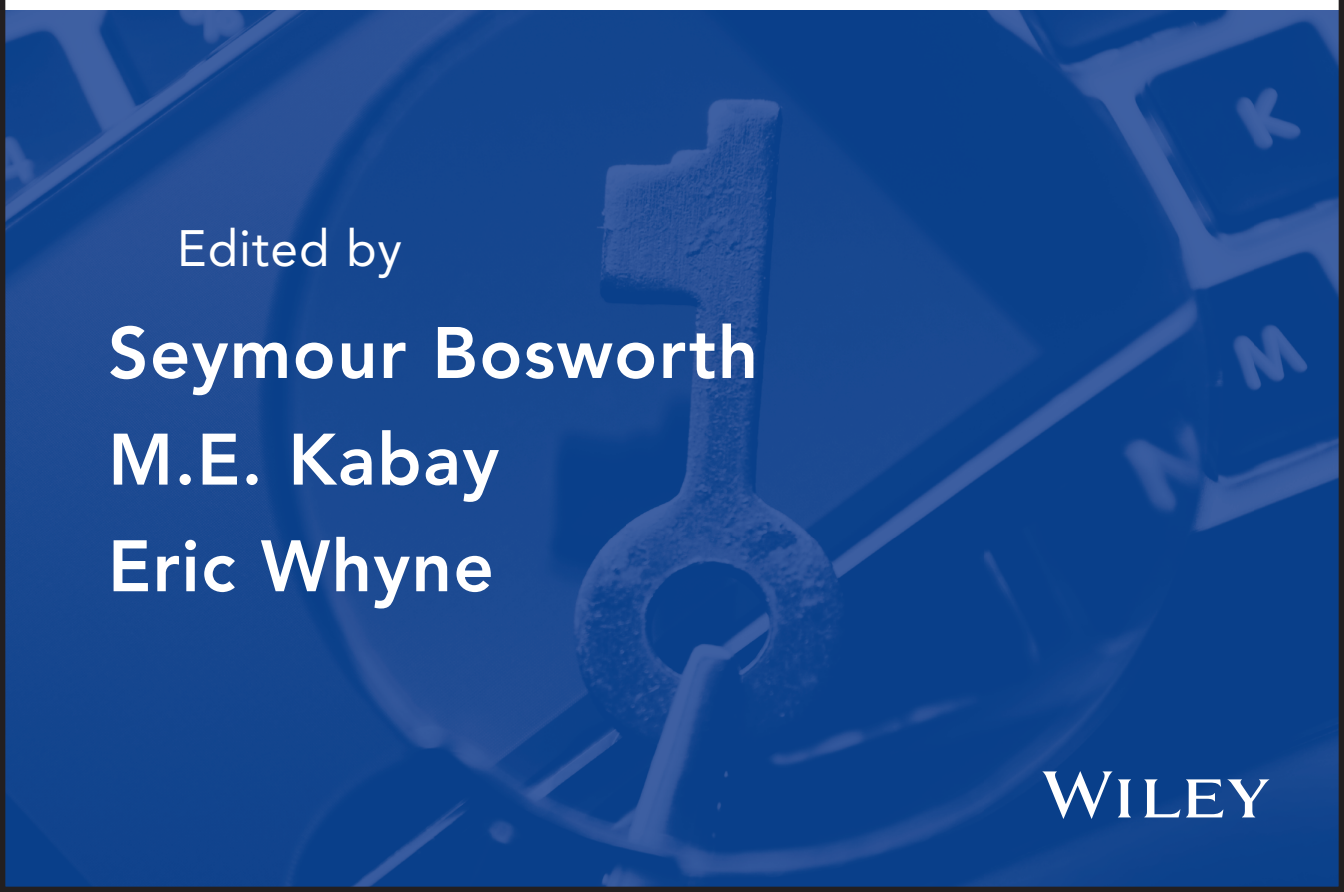
Sixth Edition

COMPUTER SECURITY HANDBOOK



Edited by

Seymour Bosworth

M.E. Kabay

Eric Whyne

WILEY

COMPUTER SECURITY HANDBOOK

COMPUTER SECURITY HANDBOOK

Sixth Edition

Volume 1

Edited by

SEYMOUR BOSWORTH

MICHEL E. KABAY

ERIC WHYNE

WILEY

Cover image: ©iStockphoto.com/Jimmy Anderson
Cover design: Wiley

Copyright © 2014 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Previous Edition: Computer Security Handbook, Fifth Edition. Copyright © 2009 by John Wiley & Sons, Inc. All Rights Reserved. Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permissions.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at http://booksupport.wiley.com. For more information about Wiley products, visit www.wiley.com.

Library of Congress Cataloging-in-Publication Data

Computer security handbook / [edited by] Seymour Bosworth, Michel E. Kabay, Eric Whyne. – Sixth edition.
volumes cm
Includes index.
ISBN 978-1-118-13410-8 (vol. 1 : pbk.) – ISBN 978-1-118-13411-5 (vol. 2 : pbk.) – ISBN 978-1-118-12706-3 (2 volume set : pbk.); ISBN 978-1-118-85174-6 (ebk); ISBN 978-1-118-85179-1 (ebk) 1. Electronic data processing departments–Security measures. I. Bosworth, Seymour. II. Kabay, Michel E. III. Whyne, Eric, 1981–
HF5548.37.C64 2014
658.4′78–dc23

2013041083

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

CONTENTS

PREFACE

ACKNOWLEDGMENTS

ABOUT THE EDITORS

ABOUT THE CONTRIBUTORS

A NOTE TO THE INSTRUCTOR

PART I FOUNDATIONS OF COMPUTER SECURITY

1. **Brief History and Mission of Information System Security**
   Seymour Bosworth and Robert V. Jacobson

2. **History of Computer Crime**
   M. E. Kabay

3. **Toward a New Framework for Information Security**
   Donn B. Parker, CISSP

4. **Hardware Elements of Security**
   Sy Bosworth and Stephen Cobb

5. **Data Communications and Information Security**
   Raymond Panko and Eric Fisher

6. **Local Area Network Topologies, Protocols, and Design**
   Gary C. Kessler

7. **Encryption**
   Stephen Cobb and Corinne LeFrançois

8. **Using a Common Language for Computer Security Incident Information**
   John D. Howard

9. **Mathematical Models of Computer Security**
Matt Bishop

10. **Understanding Studies and Surveys of Computer Crime**
M. E. Kabay

11. **Fundamentals of Intellectual Property Law**
William A. Zucker and Scott J. Nathan

PART II THREATS AND VULNERABILITIES

12. **The Psychology of Computer Criminals**
Q. Campbell and David M. Kennedy

13. **The Insider Threat**
Gary L. Tagg, CISSP

14. **Information Warfare**
Seymour Bosworth

15. **Penetrating Computer Systems and Networks**
Chey Cobb, Stephen Cobb, M. E. Kabay, and Tim Crothers

16. **Malicious Code**
Robert Guess and Eric Salveggio

17. **Mobile Code**
Robert Gezelter

18. **Denial-of-Service Attacks**
Gary C. Kessler

19. **Social-Engineering and Low-Tech Attacks**
Karthik Raman, Susan Baumes, Kevin Beets, and Carl Ness

20. **Spam, Phishing, and Trojans: Attacks Meant to Fool**
Stephen Cobb

21. **Web-Based Vulnerabilities**
Anup K. Ghosh, Kurt Baumgarten, Jennifer Hadley, and Steven Lovaas

22. **Physical Threats to the Information Infrastructure**
Franklin Platt

PART III PREVENTION: TECHNICAL DEFENSES

23. **Protecting the Physical Information Infrastructure**
Franklin Platt

24. Operating System Security
William Stallings

25. Local Area Networks
N. Todd Pritsky, Joseph R. Bumblis, and Gary C. Kessler

26. Gateway Security Devices
Justin Opatrny

27. Intrusion Detection and Intrusion Prevention Devices
Rebecca Gurley Bace

28. Identification and Authentication
Ravi Sandhu, Jennifer Hadley, Steven Lovaas, and Nicholas Takacs

29. Biometric Authentication
Eric Salveggio, Steven Lovaas, David R. Lease, and Robert Guess

30. E-Commerce and Web Server Safeguards
Robert Gezelter

31. Web Monitoring and Content Filtering
Steven Lovaas

32. Virtual Private Networks and Secure Remote Access
Justin Opatrny and Carl Ness

33. 802.11 Wireless LAN Security
Gary L. Tagg, CISSP and Jason Sinchak, CISSP

34. Securing VoIP
Christopher Dantos and John Mason

35. Securing P2P, IM, SMS, and Collaboration Tools
Carl Ness

36. Securing Stored Data
David J. Johnson, Nicholas Takacs, Jennifer Hadley, and M. E. Kabay

37. PKI and Certificate Authorities
Santosh Chokhani, Padgett Peterson, and Steven Lovaas

38. Writing Secure Code
Lester E. Nichols, M. E. Kabay, and Timothy Braithwaite

39. Software Development and Quality Assurance
Diane E. Levine, John Mason, and Jennifer Hadley

40. Managing Software Patches and Vulnerabilities
Karen Scarfone, Peter Mell, and Murugiah Souppaya

41. Antivirus Technology
Chey Cobb and Allysa Myers

42. Protecting Digital Rights: Technical Approaches
Robert Guess, Jennifer Hadley, Steven Lovaas, and Diane E. Levine

PART IV PREVENTION: HUMAN FACTORS

43. Ethical Decision Making and High Technology
James Landon Linderman

44. Security Policy Guidelines
M. E. Kabay and Bridgitt Robertson

45. Employment Practices and Policies
M. E. Kabay and Bridgitt Robertson

46. Vulnerability Assessment
Rebecca Gurley Bace and Jason Sinchak

47. Operations Security and Production Controls
M. E. Kabay, Don Holden, and Myles Walsh

48. Email and Internet Use Policies
M. E. Kabay and Nicholas Takacs

49. Implementing a Security-Awareness Program
K. Rudolph

50. Using Social Psychology to Implement Security Policies
M. E. Kabay, Bridgitt Robertson, Mani Akella, and D. T. Lang

51. Security Standards for Products
Paul Brusil and Noel Zakin

PART V DETECTING SECURITY BREACHES

52. Application Controls
Myles Walsh and Susan Baumes

53. Monitoring and Control Systems
Caleb S. Coggins and Diane E. Levine

54. Security Audits
Donald Glass, Richard O. Moore III, Chris Davis, John Mason, David Gursky, James Thomas, Wendy Carr, M. E. Kabay, and Diane Levine

55. Cyber Investigation
Peter Stephenson

PART VI RESPONSE AND REMEDIATION

56. Computer Security Incident Response Teams
Michael Miora, M. E. Kabay, and Bernie Cowens

57. Data Backups and Archives
M. E. Kabay and Don Holden

58. Business Continuity Planning
Michael Miora

59. Disaster Recovery
Michael Miora

60. Insurance Relief
Robert A. Parisi, Jr., John F. Mullen, and Kevin Apollo

61. Working with Law Enforcement
David A. Land

PART VII MANAGEMENT'S ROLE IN SECURITY

62. Quantitative Risk Assessment and Risk Management
Robert V. Jacobson and Susan Baumes

63. Management Responsibilities and Liabilities
Carl Hallberg, M. E. Kabay, Bridgitt Robertson, and Arthur E. Hutt

64. U.S. Legal and Regulatory Security Issues
Timothy Virtue

65. The Role of the CISO
Karen F. Worstell

66. Developing Security Policies
M. E. Kabay and Sean Kelley

67. Developing Classification Policies for Data
Karthik Raman, Kevin Beets, and M. E. Kabay

68. Outsourcing and Security
Kip Boyle, Michael Buglewicz, and Steven Lovaas

PART VIII PUBLIC POLICY AND OTHER CONSIDERATIONS

69. Privacy in Cyberspace: U.S. and European Perspectives
Henry L. Judy, Scott L. David, Benjamin S. Hayes, Jeffrey B. Ritter, Marc Rotenberg, and M. E. Kabay

70. Anonymity and Identity in Cyberspace
M. E. Kabay, Eric Salveggio, Robert Guess, and Russell D. Rosco

71. Healthcare Security and Privacy
Paul Brusil

72. Legal and Policy Issues of Censorship and Content Filtering
Lee Tien, Seth Finkelstein, and Steven Lovaas

73. Expert Witnesses and the *Daubert* Challenge
Chey Cobb

74. Professional Certification and Training in Information Assurance
M. E. Kabay, Christopher Christian, Kevin Henry, and Sondra Schneider

75. The Future of Information Assurance
Jeremy A. Hansen

PREFACE

Computers are an integral part of our economic, social, professional, governmental, and military infrastructures. They have become necessities in virtually every area of modern life, but their vulnerability is of increasing concern. Computer-based systems are constantly under threats of inadvertent error and acts of nature, as well as those attributable to unethical, immoral, and criminal activities. It is the purpose of *The Computer Security Handbook* to provide guidance in recognizing these threats, eliminating them where possible and, if not, then reducing any losses attributable to them.

The Handbook will be most valuable to those directly responsible for computer, network, or information security, as well as those who must design, install, and maintain secure systems. It will be equally important to those managers whose operating functions can be affected by breaches in security and to those executives who are responsible for protecting the assets that have been entrusted to them.

With the advent of desktop, laptop, and handheld computers, and with the vast international networks that interconnect them, the nature and extent of threats to computer security have grown almost beyond measure. In order to encompass this unprecedented expansion, *The Computer Security Handbook* has grown apace.

When the first edition of the *Handbook* was published, its entire focus was on mainframe computers, the only type then in widespread use. The second edition recognized the advent of small computers, while the third edition placed increased emphasis on PCs and networks.

Edition	Publication Date	Chapters	Text Pages
First	1973	12	162
Second	1988	19	383
Third	1995	23	571
Fourth	2002	54	1,184
Fifth	2009	77	2,040
Sixth	2014	75	2,224

The fourth edition of *The Computer Security Handbook* gave almost equal attention to mainframes and microcomputers, requiring more than twice the number of chapters and pages as the third.

The fifth edition was as great a step forward as the fourth. With 77 chapters and the work of 86 authors, we increased coverage in both breadth and depth. In this sixth edition, we updated all chapters while continuing to cover all 10 domains of the Common Body of Knowledge, as defined by the International Information Systems Security Certification Consortium (ISC)[2]:

1. Security Management Practices: Chapters 10, 12, 13, 14, 15, 19, 10, 31, 43, 44, 45, 46, 47, 48, 49, 50, 51, 54, 55, 62, 63, 64, 65, 66, 67, 68, 74, 75
2. Security Architecture and Models: Chapters 1, 2, 3, 8, 9, 24, 26, 27, 51
3. Access Control Systems and Methodology: Chapters 15, 19, 28, 29, 32
4. Application Development Security: Chapters 13, 19, 21, 30, 38, 39, 52, 53
5. Operations Security: Chapters 13, 14, 15, 19, 21, 24, 36, 40, 47, 53, 57
6. Physical Security: Chapters 4, 13, 15, 19, 22, 23, 28, 29
7. Cryptography: Chapters 7, 32, 37, 42
8. Telecomm, Networks, and Internet Security: Chapters 4, 5, 6, 13, 14, 15, 16, 17, 18, 20, 21, 24, 25, 26, 27, 30, 31, 32, 33, 34, 35, 41, 48
9. Business Continuity Planning: Chapters 22, 23, 56, 57, 58, 59, 60
10. Law, Investigations, and Ethics: Chapters 11, 12, 13, 31, 42, 61

We have continued our practice from the fourth and fifth editions of inviting a security luminary to write the final chapter, "The Future of Information Assurance." We are pleased to include this stellar contribution from Jeremy A. Hansen.

SEYMOUR BOSWORTH
Editor-in-Chief
February 2014

ACKNOWLEDGMENTS

Seymour Bosworth, Editor-in-Chief. I would like to give grateful recognition to Arthur Hutt and Douglas Hoyt, my coeditors of the first, second, and third editions of this *Handbook.* Although both Art and Doug are deceased, their commitment and their competence remain as constant reminders that nothing less than excellence is acceptable. Mich Kabay, my coeditor from the fourth and fifth editions, and Eric Whyne, our fellow editor from the fifth and now sixth editions, continue in that tradition. I would not have wanted to undertake this project without them.

Thanks are also due to our colleagues at John Wiley & Sons: Tim Burgard as former Acquisitions Editor, Helen Cho as Editorial Program Coordinator, Sheck Cho as Executive Editor, Kimberly Kappmeyer as Production Editor, Natasha Andrews as Senior Production Editor, and Darice Moore as Copyeditor. All have performed their duties in an exemplary manner and with unfailing kindness, courtesy, and professionalism.

M. E. Kabay, Technical Editor. I want to thank my beloved wife, Deborah Black, light of my life, for her support and understanding over the years that this project has taken away from our time together. I am also grateful to the authors who have selflessly contributed so much to updating the material presented in this text.

Eric Whyne, Administrative Editor. An undertaking as big as pulling together this handbook would not be possible without my wife Lindsay and the love and support she gives to me and to our son Colton. I'd also like to thank the friends and mentors that have helped me most in my career: Mich and Sy, Tom Aldrich, Tom Payne, Frank Vanecek, and my parents Len and Terri. Any successful undertakings I've had, including this book, have been from listening to the advice they've given and aspiring to internalize the virtues that they exemplify. The authors who have contributed to this book also deserve many thanks for sharing their experience and wisdom. It is something for which I, myself, and the readers are extremely grateful.

ABOUT THE EDITORS

Seymour Bosworth, M.S., CDP (email: sybosworth55@gmail.com) is president of S. Bosworth & Associates, Plainview, New York, a management consulting firm specializing in computing applications for banking, commerce, and industry. Since 1972, he has been a contributing editor for all six editions of the *Computer Security Handbook,* and for several editions has been Editor-in-Chief. He has written many articles and lectured extensively about computer security and other technical and managerial subjects. He has been responsible for the design and manufacture, systems analysis, programming, and operations, of both digital and analog computers. For his technical contributions, including an error-computing calibrator, a programming aid, and an analog-to-digital converter, he has been granted a number of patents, and is working on several others.

Bosworth is a former president and CEO of Computer Corporation of America, manufacturers of computers for scientific and engineering applications; president of Abbey Electronics Corporation, manufacturers of precision electronic instruments and digital devices; and president of Alpha Data Processing Corporation, a general-purpose computer service bureau. As a vice president at Bankers Trust Company, he had overall responsibility for computer operations, including security concerns.

For more than 20 years, Bosworth was an adjunct associate professor of management at the Information Technologies Institute of New York University, where he lectured on computer security and related disciplines. He has conducted many seminars and training sessions for the Battelle Institute, New York University, the Negotiation Institute, the American Management Association, and other prestigious organizations. For many years he served as arbitrator, chief arbitrator, and panelist for the American Arbitration Association. He holds a master's degree from the Graduate School of Business of Columbia University and a Certificate in Data Processing from the Data Processing Management Association.

M. E. Kabay, Ph.D., CISSP-ISSMP (email: mekabay@gmail.com) has been programming since 1966. In 1976, he received his Ph.D. from Dartmouth College in applied statistics and invertebrate zoology. After joining a compiler and relational database team in 1979, he worked for Hewlett-Packard (Canada) Ltd. from 1980 through 1983 as an HP3000 operating system performance specialist and then ran operations at a large service bureau in Montréal in the mid-1980s before founding his own operations management consultancy. From 1986 to 1996, he was an adjunct instructor in the John Abbott College professional programs in programming and in technical support. He was director of education for the National Computer Security Association from 1991 to the end of 1999 and was security leader for the INFOSEC Group of AtomicTangerine, Inc., from January 2000 to June 2001. In July 2001, he joined the

faculty at Norwich University as associate professor of computer information systems in the School of Business and Management. In January 2002, he took on additional duties as the director of the graduate program in information assurance in the School of Graduate and Continuing Studies at Norwich, where he was also chief technical officer for several years. He returned to full-time teaching in the School of Business and Management in 2009 and was promoted to professor of computer information systems in 2011. He serves as associate director of the Norwich University Center for Advanced Computing and Digital Forensics.

Kabay was inducted into the Information Systems Security Association Hall of Fame in 2004. He has published more than 1,500 articles in operations management and security in several trade journals since 1986. He wrote two columns a week for *Network World Security Strategies* between 2000 and 2011; archives are at www.mekabay.com/nwss. For the last three editions, Kabay has been Technical Editor of the *Computer Security Handbook.* He also has a Website with freely available teaching materials and papers at www.mekabay.com.

Eric Whyne (email: ericwhyne@gmail.com), administrative editor of the *Computer Security Handbook*, is a technical manager and engineer at Data Tactics Corporation where he develops solutions that benefit national security, currently managing and working on several DARPA-funded big data and data science projects. Prior to this position, he was employed by Exelis Corp. and managed the Joint Improvised Explosive Device Defeat Organization (JIEDDO) Counter-IED Operations Integration Center (COIC) Systems Integration Laboratory (SIL), which consisted of engineers and analysts tasked to develop, deploy, and maintain global software and systems designed to provide intelligence to help predict and prevent explosive devices in Iraq and Afghanistan. Previously, he worked as an engineer with Pennsylvania State University Applied Research Labs (PSU ARL) researching the development and use of immersive visualization systems, geospatial information systems, visual data mining, and deploying touch interfaces in operational environments.

Prior to his industry experience, Whyne spent nine years on active duty in the United States Marine Corps (USMC) in ground combat units, starting as enlisted and attaining the rank of captain. His accomplishments in the Marine Corps include two meritorious promotions and a Navy Commendation Medal with Valor distinction for actions in combat. During his time in the military, he worked in the fields of signals intelligence and communications and served as an advisor to the Iraqi Army. Since 2005, he has been the coordinating editor for the 5th and 6th editions of the *Computer Security Handbook*. He contributes to several open source projects, serves as an invited technical expert in the W3C HTML5 working group, and is a member of the AFCEA Technology Committee.

Whyne attended Norwich University and graduated magna cum laude with a B.S. in computer science and minor degrees in mathematics, information assurance, and engineering.

ABOUT THE CONTRIBUTORS

Wendy Adams Carr currently works for the U.S. Army Corps of Engineers as a member of the Computer Incident Response Team (CIRT). Prior to this she performed as an Information Assurance Security Engineer with Booz Allen & Hamilton, where she supported a Department of Defense client in developing and maintaining DITSCAP and DIACAP-based certification and accreditation of complex, large-scale Information Systems. She is retired from the U.S. Army. She is also an active member of Infragard.

Mani Akella, a director (technology), has been actively working with information-security architectures and identity protection for Consultantgurus and its clients. An industry professional for 20 years, he has worked with hardware, software, networking, and all the associated technologies that service information in all of its incarnations and aspects. Over the years, he has developed a particular affinity for international data law and understanding people and why they do what they do (or do not). He firmly believes that the best law and policy is that which understands and accounts for cross-cultural differences, and works with an understanding of culture and societal influences. To that end, he has been actively working with all his clients and business acquaintances to improve security policies and make them more people-friendly: His experience has been that the best policy is that which works with, instead of being antagonistic to, the end user.

Rebecca Gurley Bace is the president/CEO of Infidel, Inc., a strategic consulting practice headquartered in Scotts Valley, California. She is also a venture consultant for Palo Alto–based Trident Capital, where she is credited with building Trident's investment portfolio of security product and service firms. Her areas of expertise include intrusion detection and prevention, vulnerability analysis and mitigation, and the technical transfer of information-security research results to the commercial product realm. Prior to transitioning to the commercial world, she worked in the public sector, first at the National Security Agency, where she led the Intrusion Detection research program, then at the Computing Division of the Los Alamos National Laboratory, where she served as deputy security officer. Her publishing credits include two books, an NIST Special Publication on intrusion detection and prevention, and numerous articles on information-security technology topics.

Susan Baumes, MS, CISSP, is an information-security professional working in the financial services industry. In her current role, she works across the enterprise to develop information-security awareness and is responsible for application security. Her role also extends to policy development, compliance, and audit. She has 11 years'

experience in application development, systems and network administration, database management, and information security. Previously, she worked in a number of different sectors, including government (federal and state), academia, and retail.

Kurt Baumgarten, CISA, is vice president of information security and a partner at Peritus Security Partners, LLC, a leader in providing compliance-driven information security solutions. He is also a lecturer, consultant, and the developer of the DDIPS intrusion prevention technology as well as a pioneer in using best practices frameworks for the improvement of information technology security programs and management systems. He has authored multiple articles about the business benefits of sound information technology and information assurance practices, and assists businesses and government agencies in defining strategic plans that enhance IT and IA as positive value chain modifiers. He holds both a master's of science in information assurance and an M.B.A. with a concentration in e-commerce, and serves as an adjunct professor of information assurance. He has more than 20 years of experience in IT infrastructure and information security and is an active member of ISSA, ISACA, ISSSP, and the MIT Enterprise Forum. He periodically acts as an interim Director within external organizations in order to facilitate strategic operational changes in IT and information security.

Kevin Beets has been a research scientist with McAfee for over nine years. His work has concentrated on vulnerability, exploit and malware analysis, and documentation for the Foundstone and McAfee Labs teams. Prior to working with McAfee, he architected private LANS as well as built, monitored, and supported CheckPoint and PIX firewalls and RealSecure IDS systems.

Matt Bishop is a professor in the Department of Computer Science at the University of California at Davis and a codirector of the Computer Security Laboratory. His main research area is the analysis of vulnerabilities in computer systems, especially their origin, detection, and remediation. He also studies network security, policy modeling, and electronic voting. His textbook, *Computer Security: Art and Science*, is used widely in advanced undergraduate and graduate courses. He received his Ph.D. in computer science from Purdue University, where he specialized in computer security, in 1984.

Kip Boyle is the chief information-security officer of PEMCO Insurance, a $350 million property, casualty, and life insurance company serving the Pacific Northwest. Prior to joining PEMCO Insurance, he held such positions as chief security officer for a $50 million national credit card transaction processor and technology service provider; authentication and encryption product manager for Cable & Wireless America; senior security architect for Digital Island, Inc.; and a senior consultant in the Information Security Group at Stanford Research Institute (SRI) Consulting. He has also held director-level positions in information systems and network security for the U.S. Air Force. He is a Certified Information System Security Professional and Certified Information Security Manager. He holds a bachelor's of science in computer information systems from the University of Tampa (where he was an Air Force ROTC Distinguished Graduate) and a master's of science in management from Troy State University.

Jennifer Bradley is a member of the first Master of Science in Information Assurance graduating class at Norwich University. She is the primary Systems and Security Consultant for Indiana Networking in Lafayette, Indiana, and has served as both a

network and systems administrator in higher education and private consulting. She has almost 15 years' experience as a programmer and instructor of Web technologies, with additional interests in data backup, virtualization, authentication/identification, monitoring, desktop and server deployment, and incident response. At present she serves as an independent consultant. She has previously worked as a tester for quality and performance projects for Google, Inc., and as a collegiate adjunct instructor in computer technologies. She received a bachelor's of science in Industrial and Computer Technology from Purdue University.

Timothy Braithwaite has more than 30 years of hands-on experience in all aspects of automated information processing and communications. He is currently the deputy director of strategic programs at the Center for Information Assurance of Titan Corporation. Before joining Titan, he managed most aspects of information technology, including data and communications centers, software development projects, strategic planning and budget organizations, system security programs, and quality improvement initiatives. His pioneering work in computer systems and communications security while with the Department of Defense resulted in his selection to be the first systems security officer for the Social Security Administration (SSA) in 1980. After developing security policy and establishing a nationwide network of regional security officers, he directed the risk assessment of all payment systems for the agency. In 1982, he assumed the duties of deputy director, systems planning and control of the SSA, where he performed substantive reviews of all major acquisitions for the associate commissioner for systems and, through a facilitation process, personally led the development of the first Strategic Systems Plan for the administration. In 1984, he became director of information and communication services for the Bureau of Alcohol, Tobacco, and Firearms at the Department of Treasury. In the private sector, he worked in senior technical and business development positions for SAGE Federal Systems, a software development company; Validity Corporation, a testing and independent validation and verification company; and J.G. Van Dyke & Associates, where he was director, Y2K testing services. He was recruited to join Titan Corporation in December 1999 to assist in establishing and growing the company's Information Assurance practice.

Dr. Paul Brusil founded Strategic Management Directions, a security and enterprise management consultancy in Beverly, Massachusetts. He has been working with various industry and government sectors, including healthcare, telecommunications, and middleware to improve the specification, implementation, and use of trustworthy, quality, security-related products and systems. He supported strategic planning that led to the National Information Assurance Partnership and other industry forums created to understand, promote, and use the Common Criteria to develop security and assurance requirements and to evaluate products. He has organized, convened, and chaired several national workshops, conferences, and international symposia pertinent to management and security. Through these and other efforts to stimulate awareness and cooperation among competing market forces, he spearheaded industry's development of the initial open, secure, convergent, standards-based network and enterprise management solutions. While at the MITRE Corp, he led research and development critical to the commercialization of the world's first LAN solutions. Earlier, at Harvard, he pioneered research leading to noninvasive diagnosis of cardiopulmonary dysfunction. He is a Senior Member of the IEEE, a member of the Editorial Advisory Board of the *Journal of Network and Systems Management* (JNSM), has been senior technical editor for JNSM, is the guest editor for all JNSM's Special Issues on Security and Management,

and is a lead instructor for the adjunct faculty supporting the master's of science in information assurance degree program at Norwich University. He has authored over 100 papers and book chapters. He graduated from Harvard University with a joint degree in Engineering and Medicine.

Michael Buglewicz is employed at National Security Technologies as a section manager whose team provides technology and communications solutions to various government agencies. He spent 17 years at Microsoft in various roles in services and business management. Prior to Microsoft, he was involved with building some of the first Internet ecommerce and banking solutions while at First Data Corporation; he also spent ten years in law enforcement. In addition to his contributions to *The Computer Security Handbook*, he was a contributing author to *The Encyclopedia of Information Assurance* and, most recently, contributed to the book *Cloud Migration* by Tobias Hollwarth.

Dr. Joseph R. Bumblis is currently a research specialist with the Institute of Electrical and Electronic Engineers (IEEE) Twin Cities (TC) Section's Phoenix Project, where he conducts research and engineering projects in the areas of sensors, signal processing, and embedded systems design. His expertise includes computer networks, embedded systems with FPGA and SoC codesign, IT systems security, and software engineering methodologies. As an Associate Professor of Computer Engineering at the University of Wisconsin-Stout (UW-Stout), he developed computer engineering curriculum and taught courses in digital design, solid-state devices, embedded systems design, and Verilog programming. Prior to joining UW-Stout, he served as an IT systems architect at BAE Systems and held several adjunct professor positions where he taught software engineering and computer networking courses.

Q. Campbell has worked in the information-security field for over six years. He specializes in information-security threat analysis and education.

Santosh Chokhani is the founder and president of CygnaCom Solutions, Inc., an Entrust company specializing in PKI. He has made numerous contributions to PKI technology and related standards, including trust models, security, and policy and revocation processing. He is the inventor of the PKI Certificate Policy and Certification Practices Statement Framework. His pioneering work in this area led to the Internet RFC that is used as the standard for CP and CPS by governments and industry throughout the world. Before starting CygnaCom, he worked for The MITRE Corporation from 1978 to 1994. At MITRE, he was senior technical manager and managed a variety of technology research, development, and engineering projects in the areas of PKI, computer security, expert systems, image processing, and computer graphics. Chokhani obtained his master's (1971) and Ph.D. (1975) in electrical engineering/computer science from Rutgers University, where he was a Louis Bevior Fellow from 1971 to 1973.

Christopher Christian is a first lieutenant and an aviator in the United States Army. He received a bachelor's degree in Computer Information Systems at Norwich University class of 2005. His primary focus of study was Information Assurance and Security. He worked as an intern for an engineering consulting company for three years. He developed cost/analysis worksheets and floor-plan layouts to maximize workspace efficiency for companies in various industries. He graduated flight school at Fort Rucker, Alabama, where he trained on the H-60 Blackhawk. He serves as a flight

platoon leader in an air assault battalion. He is currently serving in Iraq in support of Operation Iraqi Freedom 08–09.

Chey Cobb, CISSP, began her career in information security while at the National Computer Security Association (now known as TruSecure/ICSA Labs). During her tenure as the NCSA award–winning Webmaster, she realized that Web servers often created security holes in networks and became an outspoken advocate of systems security. Later, while developing secure networks for the Air Force in Florida, her work captured the attention of the U.S. intelligence agencies. She moved to Virginia and began working for the government as the senior technical security advisor on highly classified projects. Ultimately, she went on to manage the security program at an overseas site. Now semiretired, she writes books and articles on computer security and is a frequent speaker at security conferences.

Stephen Cobb, CISSP, is an independent information-security consultant and an adjunct professor of information assurance at Norwich University, Vermont. A graduate of the University of Leeds, his areas of expertise include risk assessment, computer fraud, data privacy, business continuity management, and security awareness and education. A frequent speaker and seminar leader at industry conferences around the world, he is the author of numerous books on security and privacy as well as hundreds of articles. He cofounded several security companies whose products expanded the range of security solutions available to enterprises and government agencies. As a consultant, he has advised some of the world's largest companies on how to maximize the benefits of information technology by minimizing IT risks.

Caleb S. Coggins, MSIA, GSNA, CISSP, CISA, currently works for Sylint in the area of network forensics. Prior to Sylint, he operated in an internal audit and advisory services capacity for a healthcare IT company in Tennessee that focused on revenue and payment cycle management. Previous to that, he served in IT, security, and audit functions at Bridgestone Americas and its subsidiaries for over eight years. During his time working in the Americas and West Africa, he has enjoyed collaborating with management and teammates in identifying practical and effective solutions, while reducing risk to business operations. Prior to Bridgestone, he was the information manager for a private company as well as an information-security consultant to business clients. He holds a B.A. from Willamette University and an M.S. in information assurance from Norwich University.

Bernie Cowens, CISSP, CISA, is chief information-security officer at a Fortune 500 company in the financial services industry. He is an information risk, privacy, and security expert with more than 20 years' experience in industries including defense, high technology, healthcare, financial, and Big Four professional services. He has created, trained, and led a number of computer emergency, forensic investigation, and incident response teams over the years. He has real-world experience responding to attacks, disasters, and failures resulting from a variety of sources, including malicious attackers, criminals, and foreign governments. He has served as an advisor to and a member of national-level panels charged with analyzing cybersystem threats to critical infrastructures, assessing associated risks, and recommending both technical and nontechnical mitigation policies and procedures. He holds a master's degree in management information systems along with undergraduate degrees and certificates in systems management and information processing.

Tim Crothers is an IT manager at 3M IT.

Rob Cryan is chief information-security officer for MAPFRE USA, consisting of seven property and casualty insurance companies in the United States. He has corporate responsibility for all aspects of information security and business continuity. He has worked in both the corporate and the consulting fields, managing and delivering security solutions. His current areas of focus include managing risk in cloud computing and delivering cost-effective risk management derived from larger, often cumbersome models to apply uniform controls across multiple compliance disciplines. He received his B.S. in business administration management from the University of Maine. He is currently pursuing his M.S. in information systems and technology management with a concentration in information assurance and security.

Christopher Dantos is a senior architectural specialist with Computer Science Corporation's Global Security Solutions Group. His areas of expertise include 802.11, VoIP, and Web application security. Prior to joining CSC, he spent 10 years as a security architect with Motorola Inc., including five years in the Motorola Labs Wireless Access Research Center of Excellence. He holds a master's of science degree in information assurance from Norwich University and a bachelor's of science degree in marine engineering from the Maine Maritime Academy.

Scott L. David is executive director of the Law, Technology, and Arts Group at the University of Washington School of Law.

Chris Davis, MBA, CISA, CISSP, CCNP, finds that his role as a cloud security and compliance product manager at VCE enables him to apply his past experiences to the latest in data center computing. He has trained and presented in information security, audit, forensic analysis, hardware security design, auditing, and systems engineering for government, corporate, and university requirements. He has written or contributed to nine books covering multiple security disciplines and teaches as an adjunct professor at Southern Methodist University covering graduate courses in Information Security and Risk Management (EMIS7380) and IT Controls (EMIS7382). His professional career has taken him through Texas Instruments followed by several startup and consulting roles. He holds a bachelor's degree in nuclear engineering technologies from Thomas Edison State College and a master's in business from the University of Texas McCombs School of Business at Austin. He served eight years in the U.S. Naval Submarine Fleet onboard the special projects Submarine NR-1 and the ballistic missile submarine USS Nebraska.

Seth Finkelstein is a professional programmer with degrees in Mathematics and in Physics from MIT. He cofounded the Censorware Project, an anti-censorware advocacy group. In 1998, his efforts evaluating the sites blocked by the library's Internet policy in Loudoun County, Virginia, helped the American Civil Liberties Union win a federal lawsuit challenging the policy. In 2001, he received a Pioneer of the Electronic Frontier Award from the Electronic Frontier Foundation for his groundbreaking work in analyzing content-blocking software. In 2003, he was primarily responsible for winning a temporary exemption in the Digital Millennium Copyright Act allowing for the analysis of censorware.

Eric Fisher is a systems architecture and network security engineer for Penn State University's Applied Research Laboratories and is currently a master's candidate

in the field of information security and forensics. Eric has extensive experience designing and implementing high performance and redundant compute systems for the Department of Defense and the IC at large. His areas of expertise are designing and implementing secure and scalable information systems, and the networks that connect them. Before joining Penn State, he worked for the Raytheon Corporation as a *nix/Windows/Network administrator where he managed large-scale high-availability clusters.

Robert Gezelter, CDP, has over 33 years of experience in computing, starting with programming scientific/technical problems. Shortly thereafter, his focus shifted to operating systems, networks, security, and related matters, where he has 32 years of experience in systems architecture, programming, and management. He has worked extensively in systems architecture, security, internals, and networks, ranging from high-level strategic issues to the low-level specification, design, and implementation of device protocols and embedded firmware. He is an alumnus of the IEEE Computer Society's Distinguished Visitor Program for North America, having been appointed to a three-year term in 2004. His appointment included many presentations at Computer Society chapters throughout North America. He has published numerous articles, appearing in *Hardcopy, Computer Purchasing Update, Network Computing, Open Systems Today, Digital Systems Journal,* and *Network World*. He is a frequent presenter at conference sessions on operating systems, languages, security, networks, and related topics at local, regional, national, and international conferences, speaking for DECUS, Encompass, IEEE, ISSA, ISACA, and others. He previously authored the mobile code and Internet-related chapters for the 4th edition of this *Handbook* (2002) as well as the "Internet Security" chapters of the 3rd edition (1995) and its supplement (1997). He is a graduate of New York University with B.A. (1981) and M.S. (1983) degrees in computer science. He founded his consulting practice in 1978, working with clients both locally and internationally. He maintains his offices in Flushing, New York. He may be contacted via his firm's Website at www.rlgsc.com.

Dr. Anup K. Ghosh is president and chief executive of Secure Command, LLC, a security software start-up developing next-generation Internet security products for corporate networks. He also holds a position as research professor at George Mason University. He was previously senior scientist and program Manager in the Advanced Technology Office of the Defense Advanced Research Projects Agency (DARPA), where he managed an extensive portfolio of information assurance and information operations programs. He previously served in executive management as Vice President of Research at Cigital, Inc. He has served as principal investigator on contracts from DARPA, NSA, and NIST's Advanced Technology Program and has written more than 40 peer-reviewed conference and journal articles. He is also author of three books on computer network defense, serves on the editorial board of *IEEE Security and Privacy Magazine*, and has been guest editor for *IEEE Software* and *IEEE Journal on Selected Areas in Communications*. He is a Senior Member of the IEEE. For his contributions to the Department of Defense's information assurance, he was awarded the Frank B. Rowlett Trophy for Individual Contributions by the National Security Agency in November 2005, a federal government–wide award. He was also awarded the Office of the Secretary of Defense Medal for Exceptional Public Service for his contributions while at DARPA. In 2005, Worcester Polytechnic Institute awarded him its Hobart Newell Award for Outstanding Contributions to the Electrical and Computer Engineering Profession. He has previously been awarded the IEEE's Millennium Medal for Outstanding Contributions to E-Commerce Security. He completed his Ph.D. and

master's of science in electrical engineering at the University of Virginia and his bachelor's of science in electrical engineering at Worcester Polytechnic Institute.

Donald Glass, CISA, CISSP, has over 15 years of experience in the IT Auditing and Information Security fields. He is the current director of IT audit for Kerzner International. Author of several information security and IT audit articles, he is recognized as a leader in the IT audit field and information security.

Robert Guess, CISSP, NSA-IAM, NSA-IEM, is a senior security engineer at a Fortune 500 organization and an Associate Professor of Information Systems Technology. He possesses an M.S. in information assurance and has over a dozen industry certifications, including the Certified Information Systems Security Practitioner (CISSP), National Security Agency INFOSEC Assessment Methodologist (NSA-IAM), and National Security Agency INFOSEC Evaluation Methodologist (NSA-IEM). In addition to his academic experience, his professional experience includes work within the air, space, and defense sectors, serving as primary subject matter expert on a National Science Foundation cybersecurity grant, the development of workforce certification standards for information assurance professionals (DOD 8570.01-m), and periodic work as both an editor and author (e.g., the 5th and 6th editions of *The Computer Security Handbook*, etc.). His work in recent years has focused on penetration testing, incident response, the forensic analysis of digital evidence, virtualization, and cloud computing.

David Gursky, CISA, CISM, CISSP, is an information assurance manager and researcher at Raytheon Integrated Defense Systems working in Crystal City, Virginia. He is principal investigator for behavior-based intrusion detection systems, attribute-based access control, and resource-efficient authentication techniques. He held several senior positions as a Department of Defense contractor supporting information assurance programs and has over 30 years' experience in information technology and information security. He has conducted numerous security audits for PriceWaterhouse and Coopers. He has a bachelor's of science degree in business management from Southern New Hampshire University, a master's of science degree from Norwich University, and an M.B.A. from Northeastern University. In addition, he holds a CISA, CISM, and CISSP certifications. He lives in Northern Virginia and is an active member of $(ISC)^2$ and ISACA.

Jennifer Hadley is a member of the first master's of science in information assurance graduating class at Norwich University. She is the primary systems and security consultant for Indiana Networking in Lafayette, Indiana, and has served as both a network and systems administrator in higher education and private consulting. She has almost 10 years' experience as a programmer and instructor of Web technologies with additional interests in data backup, virtualization, authentication/identification, monitoring, desktop and server deployment, and incident response. At present, she serves as a technology consultant for Axcell Technologies, Inc. Previously she worked as a tester for quality and performance projects for Google, Inc., and as a collegiate adjunct instructor in computer technologies. She received a bachelor's of science degree in industrial and computer technology from Purdue University.

Carl Hallberg, CISSP, has been a UNIX systems administrator for years as well as an information-security consultant. He has also written training courses for subjects including firewalls, VPNs, and home network security. He has a bachelor's degree in

psychology. Currently, he is a senior member of an incident response team for a major U.S. financial institution.

Jeremy A. Hansen is an information-security and cryptography educator and researcher.

David Harley, CITP, FBCS, CISSP, is CEO of Small Blue-Green World, COO of AVIEN, and ESET Senior Research Fellow, specializing in antimalware research and security/technical authoring and editing in a number of areas. His books include *Viruses Revealed* and *The AVIEN Malware Defense Guide*. Previously, he managed the UK National Health Service's Threat Assessment Centre, and before that he served as security analyst for the Imperial Cancer Research Fund (now Cancer Research UK). His academic background is in social sciences and computer science, and he is a fellow of the BCS Institute (formerly the British Computer Society).

Benjamin S. Hayes is an IT specialist at Travis Software.

Kevin Henry has been involved in computers since 1976, when he was an operator on the largest minicomputer system in Canada. He has since worked in many areas of information technology, including computer programming, systems analysis, and information technology audit. He was asked to become director of security based on the evidence of his audits and involvement in promoting secure IT operations. Following 20 years in the telecommunications field, he moved to a senior auditor position with the State of Oregon, where he was a member of the Governor's IT Security Subcommittee and performed audits on courts and court-related IT systems. He has extensive experience in risk management and business continuity and disaster recovery planning. He frequently presents papers at industry events and conferences and is on the preferred speakers list for nearly every major security conference. Since joining $(ISC)^2$ as their first full-time program manager in 2002, he has been responsible for research and development of new certifications, courseware, and development of educational programs and instructors. He has also been providing support services and consulting for organizations that require in-depth risk analysis and assistance with specific security-related challenges. This has led to numerous consulting engagements in the Middle East and Asia for large telecommunications firms, government departments, and commercial enterprises.

Don Holden, CISSP, ISSMP, is a principal consultant with Concordant specializing in information security. He has more than 20 years of management experience in information systems, security, encryption, business continuity, and disaster recovery planning in both industry and government. Previously he was a technology leader for RedSiren Technologies (formerly SRI Consulting). His achievements include leading a cyberinsurance joint venture project, developing privacy and encryption policies for major financial institutions, and recommending standards-based information technology security policies for a federal financial regulator. Holden is an adjunct professor for the Norwich University's master's of science in information assurance. He received an M.B.A. from Wharton and is a Certified Information System Security Professional and Information System Security Management Professional.

Dr. John D. Howard is a former Air Force engineer and test pilot who currently works in the Security and Networking Research Group at the Sandia National Laboratories, Livermore, California. His projects include development of the SecureLink software for

automatic encryption of network connections. He has extensive experience in systems development, including an aircraft–ground collision avoidance system for which he holds a patent. He is a graduate of the Air Force Academy, has master's degrees in both aeronautical engineering and political science, and has a Ph.D. in engineering and public policy from Carnegie Mellon University.

Arthur E. Hutt, CCEP. The late Arthur E. Hutt was an information systems consultant with extensive experience in banking, industry, and government. He served as a contributing editor to the 1st, 2nd, and 3rd editions of *The Computer Security Handbook*. He was a principal of PAGE Assured Systems, Inc., a consulting group specializing in security and control of information systems and contingency/disaster recovery planning. He was a senior information systems executive for several major banks active in domestic and international banking. His innovative and pioneering development of online banking systems received international recognition. He was also noted for his contributions to computer security and to information systems planning for municipal government. He was on the faculty of the City University of New York and served as a consultant to CUNY on curriculum and on data processing management. He also served on the mayor's technical advisory panel for the City of New York. He was active in development of national and international technical standards, via ANSI and ISO, for the banking industry.

John B. Ippolito, CISSP, PMP, is director of information assurance services for the Allied Technology Group, Inc. He has more than 35 years of IT experience including "hands-on" experience with legacy mainframes, PCs, and virtual systems. He has supported government-wide standards efforts and coauthored NIST IT security training guidelines. He has a broad range of IT experience and is currently supporting federal agencies in the development of information assurance and privacy programs and the secure introduction of new technologies such as cloud-based systems and "bring your own device" into the business environment.

Robert V. Jacobson, CPP, CISSP. The late Robert V. Jacobson was the president of International Security Technology, Inc., a New York City–based risk-management consulting firm. He founded IST in 1978 to develop and apply superior risk-management systems. Current and past government and industry clients are located in the United States, Europe, Africa, Asia, and the Middle East. He pioneered many of the basic computer security concepts now in general use. He served as the first information system security officer at Chemical Bank, now known as J P Morgan Chase. He was a frequent lecturer and had written numerous technical articles. He held B.S. and M.S. degrees from Yale University, and was a Certified Information Systems Security Professional. He was also a Certified Protection Professional of the American Society for Industrial Security. He was a member of the National Fire Protection Association and the Information Systems Security Association. In 1991, he received the Fitzgerald Memorial Award for Excellence in Security from the New York Chapter of the ISSA.

David J. Johnson is an information-security architect with a Fortune 500 financial services company where he provides enterprise guidance and solutions on topics such as cloud computing security, data protection technologies, and biometric authentication. His prior experience includes information-security analysis, advising business and IT management and staff on security risk, and development of policies and

procedures for a Fortune 1000 company. He also has experience designing, developing, and maintaining electronic commerce (EC/EDI) infrastructure and data transfers for a national financial services company. He holds a B.S. in business administration from Oregon State University, an M.S. in information assurance from Norwich University, and multiple information-security certifications from $(ISC)^2$, ISACA, and the Cloud Security Alliance.

Henry L. Judy is of counsel to Kirkpatrick & Lockhart, a national U.S. law firm. He advises clients on a wide range of corporate and financial law matters as well as federal and state securities, legislative and regulatory matters, with a particular emphasis on financial institutions, housing finance, and technology law. He is recognized for his work on the jurisdiction and dispute resolution issues of electronic commerce. He is a graduate of Georgetown University (J.D. and A.B.).

Sean Kelley is the vice president of technology services at Evolvent Technologies in Herndon, Virginia. He is a retired Naval Officer with over 15 years of information assurance experience. Sean Kelley has received an M.A. in computer resource and information management from Webster University and an M.S. in information technology management (information assurance core competency) from the Naval Postgraduate School (NPS) in Monterey, California. He concentrated his studies on computer and network security by taking classes through the NPS Center for INFOSEC Studies and Research (NPS CISR), the world's foremost center for military research and education in information assurance (IA), defensive information warfare, and computer and network security. During his tenure at NPS, he received several certifications from the Committee of National Security Systems (CNSS), which operates under NSA. He is a Certified Information Systems Security Professional (CISSP), and is also a Project Management Professional (PMP). He has been responsible for information systems at several high-level offices in Washington, DC, and in the operational setting.

David M. Kennedy, CISSP, is TruSecure Corporation's chief of research. He directs the Research Group to provide expert services to TruSecure Corporation members, clients, and staff. He supervises the Information Security Reconnaissance (IS/R) team, which collects security-relevant information, both above- and underground in TruSecure Corporation's IS/R data collection. IS/R provides biweekly and special topic reports to IS/R subscribers. He is a retired U.S. Army Military Police officer. In his last tour of duty, he was responsible for enterprise security of five LANs with Internet access and over 3,000 personal computers and workstations. He holds a B.S. in forensic science.

Dr. Gary C. Kessler, CISSP, CCE, is an associate professor of Homeland Security at Embry-Riddle Aeronautical University in Daytona Beach, Florida, specializing in cybersecurity. He is a member of the North Florida Internet Crimes against Children (ICAC) Task Force and an adjunct faculty member at Edith Cowan University in Perth, Western Australia. From 2011 to 2012, he was the program director of the M.S. in Information Assurance program at Norwich University in Northfield, Vermont; from 2000 to 2010, he was an associate professor at Champlain College in Burlington, Vermont, where he designed and directed undergraduate and graduate programs related to information security and digital forensics. He is a Certified Information Systems Security Professional (CISSP), Certified Computer Examiner (CCE), and on the board of directors of the Consortium of Digital Forensic Specialists (CDFS). He holds a B.A.

in mathematics, an M.S. in computer science, and a Ph.D. in computing technology in education. He is the coauthor of two professional texts and over 70 articles and papers, a frequent speaker at industry events, and immediate past editor-in-chief of the *Journal of Digital Forensics, Security, and Law*. More information about him can be found at www.garykessler.net.

Dr. Minjeong Kim is associate professor of Journalism and Technical Communication at the Colorado State University. Her areas of expertise include communication law and policy, copyright law, and digital media. Her research has been published in scholarly journals, including *Communication Law and Policy*, the *Journal of the Copyright Society of the U.S.A.*, the *Journal of Computer-Mediated Communication*, and *Telecommunications Policy*. She completed her M.A. and Ph.D. from the School of Journalism and Mass Communication at the University of North Carolina at Chapel Hill. Before joining the CSU faculty in Fall 2008, she was an assistant professor at Hawaii Pacific University in Honolulu.

David A. Land served in the U.S. Army as a Counterintelligence Special Agent. With David Christie, he developed and hosted the first Department of Defense Computer Crimes Conference. Since then he has investigated espionage cases for both the Army and the Department of Energy. He serves as the information technology coordinator for Anniston City Schools in Alabama and as an adjunct professor for Norwich University, his alma mater.

David T. Lang joined the U.S. Civil Service on August 15, 2011. He has more than 30 years of experience in technical program management, counterespionage, antiterrorism, security, training, risk management, and law enforcement in private industry and the military. Before assuming his current position as director of the DCIN-TS PMO, he was the chief of enterprise architecture and security for the DCIN-TS PMO. His industry positions included director of federal security, director of information assurance and forensics, director of digital forensics, and director of external training. His military assignments included nearly a decade as a Special Agent for the Air Force Office of Special Investigations (AFOSI) and over a decade in special weapons.

Dr. David R. Lease is the Chief Solution Architect at Computer Sciences Corporation. He has over 30 years of technical and management experience in the information technology, security, telecommunications, and consulting industries. His recent projects include a $2 billion security architecture redesign for a federal law enforcement agency and the design and implementation of a secure financial management system for an organization operating in 85 countries. He is a writer and frequent speaker at conferences for organizations in the intelligence community, Department of Defense, civilian federal agencies, as well as commercial and academic organizations. He is also a peer reviewer of technical research for the IEEE Computer Society. Additionally, he is on the faculty of Norwich University and the University of Fairfax, where he teaches graduate-level information assurance courses and supervises doctoral-level research.

Corinne Lefrançois is an information assurance analyst at the National Security Agency. She graduated from Norwich University with a bachelor's of science in business administration and accounting in 2004 and is a current student in Norwich University's master's of science in information assurance program.

Dr. Diane E. Levine, DBA, CISSP, CFE, FBCI, CPS, is an internationally sought after security expert with Strategic Systems Management Ltd, a leader in the computer security field for 25 years. She is one of the initial developers of the Certification for Information System Professionals (CISSP) who contributed to the creation, acceptance, and expansion of both the program and the written examinations. Besides teaching both the CISSP and CISA review courses for many years, as an adjunct professor, she created the popular graduate program in Computer Security at New York University, a program that became the model for computer security university programs throughout the world. A prominent consultant, author, and teacher, she has had a notable career developing and implementing risk management, business continuity, and enterprise security systems in corporate, nonprofit, and government institutions and organizations throughout the world. She is widely published in trade and academic press and contributed numerous chapters to the 3rd, 4th, and 5th editions of *The Computer Security Handbook*. A sought-after public speaker and member of technical panels, she served on the board of directors for numerous professional education and certification organizations (i.e., Information System Security Professionals[ISSP], Association of Certified Fraud Examiners [ACFE], Business Continuity Institute [BCI], Contingency Planning Exchange [CPE], and the Information Security, Auditing and Control Association [ISACA]). She has published regularly in print and online at *Information Week, Information Security, Executive Security, Internet Week, The TechWriter, EarthWeb, Planet IT, Security, Technology & Design, internet.com, Smart Computing, Labmice.techtarget.net, Auerbach Publications, Journal for the EDP Auditors*, and many other prominent publications. Both through teaching and publishing, much of her professional career has been devoted to developing and expanding the specialty of computer security and mentoring others. As of 2011, she has been presented with numerous awards for excellence in and her contribution to the computer security and business continuity fields. A former high-ranking corporate executive, she currently works with corporations, organizations, and governments throughout the world and is actively involved with law enforcement agencies in the apprehension and conviction of computer security criminals and the implementation of computer security legislation. Previously published obituaries by John Wiley & Sons were never corrected and remain in print, although they are totally untrue and misleading.

Dr. James Landon Linderman is an associate professor in the Computer Information Systems department at Bentley College, Waltham, Massachusetts, where he has taught for 30 years. He is a research fellow at Bentley's Center for Business Ethics, and past vice-chair of the Faculty Senate. A resident of Fitzwilliam, New Hampshire, he is a Permanent Deacon in the Roman Catholic Diocese of Worcester, Massachusetts, and a consultant in the area of computer-assisted academic scheduling and timetable construction.

Steven Lovaas, MSIA, CISSP, is the information technology security manager for Colorado State University, where he is pursuing a Ph.D. in public communication and technology. His areas of expertise include IT security policy and architecture; communication and teaching of complex technical concepts; and security issues in both K–12 and higher education. He has taught for the M.S. program in information assurance at Norwich University, and founded OMBRA Research Associates to provide technical and security consultation for social science research. As part of his volunteer commitment to educating the next generation of scientists and engineers, he serves as the president of the Colorado Science Olympiad.

Allysa Myers is the Director of Research for West Coast Labs. Her primary responsibilities are researching and analyzing technology and security threat trends as well as reviewing and developing test methodologies. Prior to joining West Coast Labs, she spent 10 years working in the Avert group at McAfee Security, Inc. While there, she wrote for the Avert blog and *Sage* magazine, plus several external publications. She also provided training demonstrations to new researchers within McAfee, along with other groups such as the Department of Defense and McAfee Technical Support and Anti-Spyware teams. She has been a member of various security industry groups, such as the Wildlist and the Drone Armies mailing list.

John Mason is the vice president and audit manager at PacTrust Bank, a part-time director at SSAE 16 Professionals, and an adjunct professor for Norwich University's MSIA program. He has over 20 years' combined experience in internal audit, regulatory compliance, information security, SSAE 16s/SAS 70s, enterprise risk management, investigations/loss prevention, and process reengineering. He has held senior positions such as chief internal auditor and vice president of audit and compliance in a variety of companies. At two multi-billion-dollar institutions, he was the chief information security officer (CISO) and helped establish information risk-management programs, as well as designing risk-based audit programs several years before Sarbanes-Oxley. He has routinely authored, reviewed, and researched finance control policies and procedures; performed audits for governmental agencies; and managed a full spectrum of financial, operational, SOX compliance, and data processing audits. He possesses an M.B.A., a B.A. in economics, and numerous certificates, including a CISM, CISA, CGEIT, and CFE. Currently, he resides in southern California. Aside from trying to help his students maximize their value, John's favorite tenet is "Learn from the mistakes of others, because you'll never live long enough to make them all yourself."

Peter Mell is a senior computer scientist in the computer security division at the National Institute of Standards and Technology. He is the program manager for the National Vulnerability Database as well as the Security Content Automation Protocol validation program. These programs are widely adopted within the U.S. government and used for standardizing and automating vulnerability and configuration management, measurement, and policy compliance checking. He has written the NIST publications on patching, malware, intrusion detection, common vulnerability scoring system, and the common vulnerabilities and exposures standard. His research experience includes the areas of intrusion detection systems, vulnerability scoring, and vulnerability databases.

Michael Miora has designed and assessed secure, survivable, highly robust systems for industry and government over the past 30 years. He was one of the original professionals granted the CISSP in the 1990s and the ISSMP in 2004, and was accepted as a Fellow of the Business Continuity Institute (BCI) in 2005. He founded and currently serves as president of ContingenZ Corporation (www.contingenz.com). He was educated at the University of California Los Angeles and the University of California Berkeley, earning bachelor's and master's degrees in mathematics. He is an adjunct professor at Norwich University, and a frequent speaker and consultant in the area of information assurance.

Richard O. Moore III has over 20 years of information-security experience, starting with 15 years in the Marine Corps intelligence community. After leaving the Marines

as the Regimental Intelligence Chief, he worked as a security consultant for Investor's Bank and Trust, preparing them for Sarbanes Oxley and a security review of their environment. At the end of that contract, he went to work for KPMG within their security consulting practice, where he performed security reviews, audits, and penetration testing activities. He left KPMG as an associate director in the Information Security Group managing the global penetration testing team and moved to the Royal Bank of Scotland Citizens Bank, where he created the security auditing function as the senior audit manager. He has continued with the Bank and is now the senior information security manager responsible for the Technical Services group, which consists of the Forensic, Penetration Testing, Risk Assessment, Cyber Intelligence, and Information Security Consultancy services. He has served on numerous industry boards, such as Board of Directors Massachusetts InfraGard, Norwich Editorial Board, and the ISACA Academic SubCommittee, and is a contributing author to the 5th and 6th editions of *The Computer Security Handbook*.

Scott J. Nathan, Esq., is an attorney whose practice includes litigation concerning intellectual property and technology matters, computer fraud and abuse, and environmental and insurance coverage matters involving the exchange of millions of pages of documents. In addition, he advises clients about, among other things, Internet-related risks and risk avoidance, proprietary and open source software licensing, service-level agreements, and insurance coverage. He has written and spoken extensively about such issues as online privacy, cyberspace jurisdiction, and the legal issues surrounding the use of open source software. He is admitted to practice before the United States Supreme Court, the United States Court of Appeals for the First Circuit, the Federal District Court for the District of Massachusetts, and the Courts of the Commonwealth of Massachusetts. He is a member of the American Bar Association's Litigation and Computer Litigation Committees.

Carl Ness, CISSP, is a senior security analyst for the University of Iowa. He has more than 10 years' experience in the information technology and information-security fields, mainly in the academic and healthcare sector. He is a speaker, author, and educator on information assurance, including security in the academic environment, messaging security, disaster recovery and business continuity, safe home computing, and information technology operations. He previously served as a systems administrator, network administrator, information technology director, and information-security officer. He holds an M.S. degree and also provides consulting to several security software development organizations.

Lester E. Nichols III has more than 15 years' experience in information technology, including: technology computing, cybersecurity, information assurance, and enterprise and security architecture. His experiences span the medical, nonprofit, financial, and local and federal government sectors in a variety of roles, including: application developer, network engineering, security operations manager, and information-security architect. He is a published author, contributing to major security publications such as the fifth and sixth editions of *The Computer Security Handbook*, and is a frequent contributor to the $(ISC)^2$ blog and other security-related communities and lectures. He also teaches college-level computing courses and provides computing lectures at secondary schools. Prior to joining ApplyLogic, he served in various technical, consulting, and leadership roles with both federal and commercial clients, such as the Administrative Offices of U.S. Courts, KCG, Prolific Solutions, the University of California Irvine,

and the University of Phoenix, and as a federal employee with the U.S. Department of Homeland Security, Transportation Security Administration.

Justin Opatrny currently works for a Fortune 500 company specializing in security infrastructure. He also is an adjunct instructor teaching information assurance courses and an independent consultant providing technology and security services. He holds a master's degree in information assurance from Norwich University; he holds industry certifications, including CISSP, GCFA, GSNA, and CASP; and he is an active member of the Information Systems Security Association and InfraGard.

Dr. John Orlando speaks and consults on the use of social media in disaster response, marketing, sales, business development, and training. He developed and managed the Master's of Science in Business Continuity Management program at Norwich University in Vermont, created the Social Media in Disaster Response LinkedIn Group, and has written over 20 articles on business continuity and disaster response in publications such as *Continuity Insights*, *Disaster Recovery Journal*, *Global Assurance*, *CFO*, and *The Definitive Handbook on Business Continuity Management*.

Dr. Raymond Panko is a professor of information technology management in the Shidler College of Business at the University of Hawaii. His interest in security began during lunches with Donn Parker in the 1970s at SRI International and has grown ever since. His textbook on IT security, *Corporate Computer and Network Security*, is published by Prentice-Hall. His current research focuses are security for end-user applications (especially spreadsheets), how to deal with fraud, and human and organizational controls. His main teaching focus, however, remains networking. In his networking classes and textbook, he emphasizes security throughput, pointing out the security implications of network protocols and practices.

Robert A. Parisi, Jr. is the senior vice-president and national technology, network risk and telecommunications practice leader for the FINPRO unit of Marsh USA. He has spoken at various businesses, technology, legal, and insurance forums throughout the world and has written on issues affecting professional liability, privacy, technology, telecommunications, media, intellectual property, computer security, and insurance. In 2002, he was honored by *Business Insurance* magazine as one of the Rising Stars of Insurance. Immediately prior to joining Marsh, he was the senior vice-president and chief underwriting officer of eBusiness Risk Solutions (a unit of the property and casualty companies of American International Group, Inc.). He joined the AIG group of companies in 1998 as legal counsel for its Professional Liability group and held several executive and legal positions within AIG, including the position of chief underwriting officer for professional liability and technology. While at AIG, he oversaw the creation and drafting of underwriting guidelines and policies for all lines of professional liability. Prior to joining AIG, he had been in private practice, principally as legal counsel to various Lloyds of London syndicates handling a variety of professional liability lines. He graduated cum laude from Fordham College with a B.A. in economics and received his law degree from Fordham University School of Law. He is admitted to practice in New York and the U.S. District Courts for the Eastern and Southern Districts of New York.

Donn B. Parker, CISSP, Fellow of the Association for Computing Machinery and Distinguished Fellow of the Information Systems Security Association (ISSA) is a retired (1997) senior management consultant from SRI International in Menlo Park,

California, who has specialized in information security and computer crime research for 40 of his 50 years in the computer field. He is the founder of the International Information Integrity Institute (I-4), which provides a forum for exchange of security information among 75 of the largest international enterprises. He has written numerous books, papers, articles, and reports in his specialty based on interviews of over 200 computer criminals and reviews of the security of many large enterprises. He received the ISSA 1992 Individual Outstanding Achievement Award, the 1994 National Computer Systems Security Award from the US NIST and NSA, the Aerospace Computer Security Associates 1994 Distinguished Lecturer Award, and the MIS Training Institute Infosecurity News 1996 Lifetime Achievement Award. He received the ISSA Hall of Fame award and was chosen as one of five pioneers and profiled in *Information Security Magazine* in 1999.

Padgett Peterson, P.E., CISSP, IAM/IEM, has been involved with computer security and encryption for over 40 years. Since 1979 he has been employed by different elements of a major aerospace contractor. He is also an adjunct professor in the Master's of Science in Information Assurance program at Norwich University.

Franklin Platt is the founder and president of Office Planning Services, a Wall Street consultancy for 20 years headquartered in Stark, New Hampshire, since 1990. He has worked extensively in security planning, management, and preparedness in both the private and public sectors. His academic background includes business administration and electrical engineering. He has received extensive government training in emergency management, including terrorism and weapons of mass destruction, much of which is not available to the public. He holds many security certifications and is currently vetted by the FBI and by several states. His areas of expertise include: security risk management; compliance with the latest Homeland Security procedures and other federal regulations that affect the private sector; risk identification and assessment; vulnerability analysis; cost-value studies; response planning; site security surveys and compliance auditing; briefing and training; second opinion; and due diligence.

N. Todd Pritsky is the director of e-learning courseware at Hill Associates, a telecommunications training company in Colchester, Vermont. He is a senior member of the technical staff and an instructor of online, lecture, and hands-on classes. His teaching and writing specialties include e-commerce, network security, TCP/IP, and the Internet, and he also leads courses on fast packet and network access technologies. He enjoys writing articles on network security and is a contributing author of *Telecommunications: A Beginner's Guide* (McGraw-Hill/Osborne). Previously, he managed a computer center and created multimedia training programs. He holds a B.A. in philosophy and Russian/Soviet studies from Colby College.

Karthik Raman, CISSP, is a security researcher on the Adobe Secure Software Engineering Team (ASSET) where he focuses on vulnerability analysis and technical collaboration with industry partners. Before joining Adobe, he was a research scientist at McAfee Labs where he worked on threat analysis, building automation systems, malware analysis, and developing advanced anti-malware technology. He holds an M.S. degree in computer science from the University of California Irvine and B.S. degrees in computer science and computer security from Norwich University. Karthik has spoken at Infosec Southwest, SOURCE Boston, SOURCE Seattle, LayerOne, and Kaspersky Security Analyst Summit and has delivered a Black Hat Webcast.

Bridgitt Robertson has been teaching business and technology courses for over a decade years. Her multidisciplinary approach to security awareness analyzes threats in the global enterprise and investigates how an educated workforce can mitigate risks and enhance corporate competitiveness. Prior to teaching, she worked for global companies in the areas of project management, business analysis, and consulting. She is a member of InfraGard.

Russell D. Rosco is currently program chair/professor of the Computer Information Network Technology program at Ivy Tech Community College–NorthEast. He is also a graduate student at Purdue University, where he is working on his dissertation for an Interdisciplinary Ph.D. in information assurance; combining his interests of computers, business, and psychology. He has an M.S. in information assurance from Norwich University, a B.S. in psychology, and a B.A. in business administration from Washington State University. Mr. Rosco had several years of IT work experience prior to his decade of teaching IT courses, ranging from intro to PCs to Active Directory design.

Marc Rotenberg is executive director of the Electronic Privacy Information Center in Washington, DC. He teaches information privacy law at Georgetown University Law Center. He has published many articles in legal and scientific journals. He is the coauthor of several books, including *Information Privacy Law*, *Privacy and Human Rights*, *The Privacy Law Sourcebook*, and *Litigation under the Federal Open Government Laws*. He frequently testifies before the U.S. Congress and the European Parliament on emerging privacy issues. He is a Fellow of the American Bar Foundation and the recipient of several awards, including the World Technology Award in Law.

K. Rudolph is the president and chief inspiration officer for Native Intelligence, Inc. She is a Certified Information Systems Security Professional (CISSP), a Federal IT Security Professional (FITSP-M), and holds a degree from Johns Hopkins University. Her interests include the psychology of security awareness and influence as related to learning and behavior, storytelling, and security-awareness metrics. Her publications include *System Forensics, Investigation, and Response* (2010) and the second edition of *Computer Forensics Jumpstart* (2011). She was also the technical editor for Michael Solomon's book, Security Strategies in Windows Platforms and Applications. She is a named contributor and participant in the group that created *NIST Special Publication 800-16, Information Technology Security Training Requirements: A Role-and Performance-Based Model*. In 2006, she was honored by the Federal Information Systems Security Educators' Association (FISSEA) as the Security Educator of the Year. She is currently a volunteer with $(ISC)^2$'s Safe and Secure Online program.

Ravi Sandhu is cofounder and chief scientist of SingleSignOn.Net in Reston, Virginia, and professor of information technology and engineering at George Mason University in Fairfax, Virginia. An ACM and an IEEE Fellow, he is the founding editor-in-chief of *ACM's Transactions on Information and System Security*, chairman of ACM's Special Interest Group on Security, Audit and Control, and security editor for *IEEE Internet Computing*. He has published over 140 technical papers on information security. He is a popular teacher and has lectured all over the world. He has provided high-level consulting services to numerous private and government organizations.

Eric Salveggio is an information technology security professional who enjoys teaching online courses in CMIS for Liberty University and auditing for Norwich University.

He works as a trained ISO 17799, NSTISSI 4011 and 4013 consultant for Dynetics Corporation of Huntsville, Alabama, in IT Security and Auditing, and as a Private Consultant in networking, network design, and security (wired and wireless) with 10 years experience. He previously worked as the IT director for the Birmingham, Alabama, campus of Virginia College, where he opened two start-up campuses—on ground and online—and created three accredited programs (two undergrad, one graduate level) at state and federal levels in Network and Cyber Security. While in this position, he was chosen as a nominee for the 2006 Information Security Executive Award, and enjoyed being the only educational facility recognized. He was personally awarded a plaque of recognition by the Stonesoft Corporation for the same. He is a published author and photographer, and enjoys working at times as a technical editor for Pearson Education and Thomson Publishing on cyberforensics, cybersecurity, and operating systems.

Karen Scarfone is the principal consultant for Scarfone Cybersecurity. She provides cybersecurity publication consulting services to federal agencies, specializing in security automation standards and network and system security guidelines. Karen was formerly a senior computer scientist for the National Institute of Standards and Technology (NIST) and has co-authored over 50 NIST Special Publications and Interagency Reports. She holds bachelor's and master's degrees in computer science, and she has 20 years of professional experience in the IT field. Her security domains include general security engineering and administration, wired and wireless network security, host security, incident response, intrusion detection, log management, vulnerability measurement, and security automation.

Sondra Schneider is CEO and founder of Security University (SU), an information security and information assurance training and certification company. She and SU have challenged security professionals for the past 10 years, delivering hands-on tactical security classes and certifications around the world. Starting in 2008, SU set up an exam server to meet the demand for tactical security certifications. In 2005, SU refreshed the preexisting AIS security training program to the new "SU Qualified Programs," which meet and exceed security professionals requirements for hands-on tactical security "skills" training. SU delivers the Qualified/Information Security Professional and Qualified/Information Assurance Professional certifications, which are the first of their kind that measure a candidate's tactical hands-on security skills. In 2004, she was awarded Entrepreneur of the Year for the First Annual Women of Innovation Award from the Connecticut Technology Council. In 2007, she was tech editor for the popular 2007 *CEH V5 Study Guide*, and a multiple chapter author for the 2007 *CHFI Study Guide*. She sits on three advisory boards for computer security (start-up) technology companies and is a frequent speaker at computer security and wireless industry events. She is a founding member of the NYC HTCIA and IETF chapters, works closely with $(ISC)^2$, ISSA, and ISACA chapters, and the security and wireless vendor community. She specializes in information security, intrusion detection, information assurance (PKI), wireless security, and security-awareness training.

Jason Sinchak, CISSP, is an information-security consultant for Emerging Defense, LLC and an independent researcher specializing in risk assessment and penetration testing. He possesses a broad range of experience, beginning his career in application development and systems management, to operational security monitoring, enterprise security program advisory, incident response, and penetration testing. His security

experience primarily stems from industry positions and consulting, where has had the pleasure of serving multiple Fortune 500 clients and deeply focusing on performing technical security assessments in a multitude of industries and environments.

William Stallings has made a unique contribution to understanding the broad sweep of technical developments in computer networking and computer architecture. He has authored 17 titles and, counting revised editions, has published a total of over 40 books on various aspects of these subjects. In over 20 years in the field, he has been a technical contributor, technical manager, and an executive with several high-technology firms. Currently, he is an independent consultant whose clients have included computer and networking manufacturers and customers, software development firms, and leading-edge government research institutions. He created and maintains the Computer Science Student Resource Site at www.computersciencestudent.com. He is a member of the editorial board of *Cryptologia*, a scholarly journal devoted to all aspects of cryptology.

Dr. Peter Stephenson is a writer, researcher and lecturer on information assurance and risk, information warfare and counterterrorism, and digital investigation and forensics on large-scale computer networks. He has lectured extensively on digital investigation and security and has written or contributed to 14 books and several hundred articles in major national and international trade, technical, and scientific publications. He is the associate program director in the Master's of Science in Information Assurance program at the Norwich University School of Graduate Studies, where he teaches information assurance, cybercrime and cyberlaw, and digital investigation on both the graduate and undergraduate levels. He is senior research scientist at the Norwich University Applied Research Institutes, chair of the Department of Computing, and the chief information-security officer for the university. He has lectured or delivered consulting engagements for the past 23 years in 11 countries plus the United States and has been a technologist for over 40 years. He operated a successful consulting practice for over 20 years and has worked for such companies as Siemens, Tektronix, and QinetiQ (UK). He obtained his Ph.D. in computer science at Oxford Brookes University, Oxford, England, where his research was in the structured investigation of digital incidents in complex computing environments. He holds a master's of arts degree in diplomacy with a concentration in terrorism from Norwich University. He is on the editorial advisory boards of *International Journal of Digital Evidence* and the Norwich University *Journal of Information Assurance* among several others, and serves as technology editor for *SC Magazine* and the editor in chief for Norwich University Press. He is a Fellow of the Institute for Communications, Arbitration and Forensics in the United Kingdom and is a member of Michigan InfraGard and the International Federation of Information Processing Technical Committee 11, Working Group 11.9, Digital Forensics. He serves on the steering Committee of the Michigan Electronic Crime Task Force. His research is focused on information conflict.

Gary L. Tagg, CISSP, is a highly experienced information-security professional with over 20 years working in the financial and government sectors. The organizations he has worked with include Deutsche Bank, PA Consulting group, Clearstream, Pearl Assurance, and Lloyds TSB. He has performed a wide range of security roles including risk management, consulting, security architecture, policy and standards, project management, development, testing, and auditing. He is currently a risk consultant in Deutsche Bank's IT security Governance Group.

Nicholas Takacs is an information-security professional and business systems director for a long-term care insurance company. He is also an adjunct professor of information assurance at Norwich University. He has expertise in the areas of security policy management, security awareness, business continuity planning, and execution. Prior to moving into the insurance industry, he spent several years in the public utility industry focusing on the areas of regulatory compliance, disaster recovery, and identity management.

James Thomas, CISSP, is a senior partner with Norwich Security Associates, a full-spectrum information assurance consultancy headquartered in Scotland. Thomas spends most of his professional time providing policy, process, and governance advice to large banking and financial organizations in the United Kingdom and Europe. He is a 2004 graduate of the Norwich University Master's of Science in Information Assurance program. Prior to focusing his efforts in the security space, he had a long career in information technology and broadcast engineering spanning the United Kingdom and the eastern United States.

Lee Tien, Esq., is a senior staff attorney with the Electronic Frontier Foundation in San Francisco, California. He specializes in free speech and surveillance law and has authored several law review articles. He received his undergraduate degree in psychology from Stanford University and his law degree from Boalt Hall School of Law, University of California Berkeley. He is also a former newspaper reporter.

Timothy Virtue, CISSP, CISA, CCE, CFE, CIPP/G, is an accomplished information assurance leader with a focus in strategic enterprise technology risk management, information security, data privacy, and regulatory compliance. He has extensive experience with publicly traded corporations, privately held businesses, government agencies, and nonprofit organizations of all sizes.

Myles Walsh is an adjunct professor at three colleges in New Jersey: Ramapo College, County College of Morris, and Passaic County Community. For the past 12 years, he has taught courses in Microsoft Office and Web page design. He also implements small Office applications and Websites. From 1966 until 1989, he worked his way up from programmer to director in several positions at CBS, CBS Records, and CBS News. His formal education includes an M.B.A. from the Baruch School of Business and a B.B.A. from St. John's University.

Karen F. Worstell, CISM, is cofounder and principal of W Risk Group, a consultancy serving clients across multiple sectors to define due diligence to a defensible standard of care for information protection. Her areas of expertise include incident detection and management, compliance, governance, secure data management, and risk management. She is coauthor of *Evaluating the Electronic Discovery Capabilities of Outside Law Firms: A Model Request for Information and Analysis* (BNA, 2006) and is a frequent speaker and contributor in risk management and information-security forums internationally. She participates in ISACA, IIA, and the ABA Science and Technology Section, Information Security Committee, and serves as president of the Puget Sound Chapter of the ISSA.

Noel Zakin is president of RANCO Consulting LLC. He has been an information technology/telecommunications industry executive for over 45 years. He has held

managerial positions at the Gartner Group, AT&T, the American Institute of CPAs, and Unisys. These positions involved strategic planning, market research, competitive analysis, business analysis, and education and training. His consulting assignments have ranged from the Fortune 500 to small start-ups and have involved data security, strategic planning, conference management, market research, and management of corporate operations. He has been active with ACM, IFIP, and AFIPS and currently with ISSA. He holds an M.B.A. from the Wharton School.

William A. Zucker, Esq., is a partner at McCarter & English, LLP's Boston office. He serves as a senior consultant for the Cutter Consortium on legal issues relating to information technology, outsourcing, and risk management, and is a member of the American Arbitration Association's National Technology Panel and a member of the CPR Institute's working group on technology business alliances and conflict management. He has also served on the faculty of Norwich University, where he taught the intellectual property aspects of computer security. He is a trial lawyer whose practice focuses on negotiation/litigation of business transactions, outsourcing/ebusiness and technology/intellectual property. Among his publications are: "The Legal Framework for Protecting Intellectual Property in the Field of Computing and Computer Software," written for *The Computer Security Handbook*, 4th edition, coauthored with Scott Nathan; and "Intellectual Property and Open Source: Copyright, Copyleft and Other Issues for the Community User."

A NOTE TO THE INSTRUCTOR

This two-volume text will serve the interests of practitioners and teachers of information assurance. The fourth edition and fifth editions of the *Handbook* were well received in academia; at least one-quarter of all copies were bought by university and college bookstores. The design of this sixth edition continues in the same vein and includes many updates to the material.

University professors looking for texts appropriate for a two-semester sequence of undergraduate courses in information assurance will find the *Handbook* most suitable. In my own work at Norwich University in Vermont, Volume I is the text for our *IS340 Introduction to Information Assurance* and Volume II is the basis for our *IS342 Management of Information Assurance* courses.

The text will also be useful as a resource in graduate courses. In the School of Graduate and Continuing Studies at Norwich University, both volumes have been used as required and supplementary reading for the 18-month, 36-credit Master's of Science in Information Security and Assurance program (MISA).

I will continue to create, update, and post PowerPoint lecture slides based on the chapters of the *Handbook* on my Website for free access by anyone applying them to noncommercial use (e.g., for self-study, for courses in academic institutions, and for unpaid industry training); the materials will be available in the IS340 and IS342 sections:

www.mekabay.com/courses/academic/norwich/is340/index.htm
www.mekabay.com/courses/academic/norwich/is342/index.htm

M. E. Kabay
Technical Editor
October 2013

INTRODUCTION TO PART I

FOUNDATIONS OF COMPUTER SECURITY

The foundations of computer security include answers to the superficially simple question "What is this all about?" Our first part establishes a technological and historical context for information assurance so that readers will have a broad understanding of why information assurance matters in the real world. Chapters focus on principles that will underlie the rest of the text: historical perspective on the development of our field; how to conceptualize the goals of information assurance in a well-ordered schema that can be applied universally to all information systems; computer hardware and network elements underlying technical security; history and modern developments in cryptography; and how to discuss breaches of information security using a common technical language so that information can be shared, accumulated, and analyzed.

Readers also learn or review the basics of commonly used mathematical models of information-security concepts and how to interpret survey data and, in particular, the pitfalls of self-selection in sampling about crimes. Finally, the first section of the text introduces elements of law (U.S. and international) applying to information assurance. This legal framework from a layman's viewpoint provides a basis for understanding later chapters; in particular, when examining privacy laws and management's fiduciary responsibilities.

Chapter titles and topics in Part I include:

1. **Brief History and Mission of Information System Security.** An overview focusing primarily on developments in the second half of the twentieth century and the first decade of the twenty-first century
2. **History of Computer Crime.** A review of key computer crimes and notorious computer criminals from the 1970s to the mid-2000s
3. **Toward a New Framework for Information Security.** A systematic and thorough conceptual framework and terminology for discussing the nature and goals of securing all aspects of information, not simply the classic triad of confidentiality, integrity, and availability
4. **Hardware Elements of Security.** A review of computer and network hardware underlying discussions of computer and network security

5. **Data Communications and Information Security.** Fundamental principles and terminology of data communications, and their implications for information assurance
6. **Local Area Network Topologies, Protocols, and Design.** Information assurance of the communications infrastructure
7. **Encryption.** Historical perspectives on cryptography and steganography from ancient times to today as fundamental tools in securing information
8. **Using a Common Language for Computer Security Incident Information.** An analytic framework for understanding, describing, and discussing security breaches by using a common language of well-defined terms
9. **Mathematical Models of Computer Security.** A review of the most commonly referenced mathematical models used to describe information-security functions
10. **Understanding Studies and Surveys of Computer Crime.** Scientific and statistical principles for understanding studies and surveys of computer crime
11. **Fundamentals of Intellectual Property Law.** An introductory review of cyberlaw: laws governing computer-related crime, including contracts, and intellectual property (trade secrets, copyright, patents, open-source models). Also, violations (piracy, circumvention of technological defenses), computer intrusions, and international frameworks for legal cooperation

CHAPTER 1

BRIEF HISTORY AND MISSION OF INFORMATION SYSTEM SECURITY

Seymour Bosworth and Robert V. Jacobson

1.1 INTRODUCTION TO INFORMATION SYSTEM SECURITY 1·1

1.2 EVOLUTION OF INFORMATION SYSTEMS 1·3

1.2.1 1950s: Punched-Card Systems 1·4
1.2.2 Large-Scale Computers 1·4
1.2.3 Medium-Size Computers 1·5
1.2.4 1960s: Small-Scale Computers 1·6
1.2.5 Transistors and Core Memory 1·7
1.2.6 Time Sharing 1·7
1.2.7 Real-Time, Online Systems 1·7
1.2.8 A Family of Computers 1·7
1.2.9 1970s: Microprocessors 1·8
1.2.10 The First Personal Computers 1·8
1.2.11 The First Network 1·8
1.2.12 Further Security Considerations 1·9
1.2.13 The First "Worm" 1·9
1.2.14 1980s: Productivity Enhancements 1·9
1.2.15 1980s: The Personal Computer 1·9
1.2.16 Local Area Networks 1·10
1.2.17 1990s: Interconnection 1·11
1.2.18 1990s: Total Interconnection 1·12
1.2.19 Telecommuting 1·12
1.2.20 Internet and the World Wide Web 1·12
1.2.21 Virtualization and the Cloud 1·13
1.2.22 Supervisory Control and Data Acquisition 1·13

1.3 GOVERNMENT RECOGNITION OF INFORMATION ASSURANCE 1·13

1.3.1 IA Standards 1·13
1.3.2 Computers at Risk 1·14
1.3.3 InfraGard 1·18

1.4 RECENT DEVELOPMENTS 1·19

1.5 ONGOING MISSION FOR INFORMATION SYSTEM SECURITY 1·20

1.6 NOTES 1·20

1.1 INTRODUCTION TO INFORMATION SYSTEM SECURITY. The growth of computers and of information technology has been explosive. Never before has an entirely new technology been propagated around the world with such speed and with so great a penetration of virtually every human activity. Computers have brought vast benefits to fields as diverse as human genome studies, space exploration, artificial intelligence, and a host of applications from the trivial to the most life-enhancing.

Unfortunately, there is also a dark side to computers: They are used to design and build weapons of mass destruction as well as military aircraft, nuclear submarines,

and reconnaissance space stations. The computer's role in formulating biologic and chemical weapons, and in simulating their deployment, is one of its least auspicious uses.

Of somewhat lesser concern, computers used in financial applications, such as facilitating the purchase and sales of everything from matchsticks to mansions, and transferring trillions of dollars each day in electronic funds, are irresistible to miscreants; many of them see these activities as open invitations to fraud and theft. Computer systems, and their interconnecting networks, are also prey to vandals, malicious egotists, terrorists, and an array of individuals, groups, companies, and governments intent on using them to further their own ends, with total disregard for the effects on innocent victims. Besides these intentional attacks on computer systems, there are innumerable ways in which inadvertent errors can damage or destroy a computer's ability to perform its intended functions.

Because of these security problems and because of a great many others described in this volume, the growth of information systems security has paralleled that of the computer field itself. Only by a detailed study of the potential problems, and implementation of the suggested solutions, can computers be expected to fulfill their promise, with few of the security lapses that plague less adequately protected systems. This chapter defines a few of the most important terms of information security and includes a very brief history of computers and information systems, as a prelude to the works that follow.

Security can be defined as the state of being free from danger and not exposed to damage from accidents or attack, or it can be defined as the process for achieving that desirable state. The objective of information system security[1] is to optimize the performance of an organization with respect to the risks to which it is exposed.

Risk is defined as the chance of injury, damage, or loss. Thus, risk has two elements: (1) chance—an element of uncertainty, and (2) potential loss or damage. Except for the possibility of restitution, information system security actions taken today work to reduce *future* risk losses. Because of the uncertainty about future risk losses, perfect security, which implies zero losses, would be infinitely expensive. For this reason, risk managers strive to optimize the allocation of resources by minimizing the total cost of information system security measures taken and the risk losses experienced. This optimization process is commonly referred to as *risk management*.

Risk management in this sense is a three-part process:

1. Identification of material risks
2. Selection and implementation of measures to mitigate the risks
3. Tracking and evaluating of risk losses experienced, in order to validate the first two parts of the process

The purpose of this *Handbook* is to describe information security system risks, the measures available to mitigate these risks, and techniques for managing security risks. (For a more detailed discussion of risk assessment and management, see Chapters 47 and 54.)

Risk management has been a part of business for centuries. Renaissance merchants often used several vessels simultaneously, each carrying a portion of the merchandise, so that the loss of a single ship would not result in loss of the entire lot. At almost the same time, the concept of insurance evolved, first to provide economic protection against the loss of cargo and later to provide protection against the loss of buildings

by fire. Fire insurers and municipal authorities began to require adherence to standards intended to reduce the risk of catastrophes like the Great Fire of London in 1666. The Insurance Institute was established in London one year later. With the emergence of corporations as limited liability stock companies, corporate directors have been required to use prudence and due diligence in protecting shareholders' assets. Security risks are among the threats to corporate assets that directors have an obligation to address.

Double-entry bookkeeping, another Renaissance invention, proved to be an excellent tool for measuring and controlling corporate assets. One objective was to make insider fraud more difficult to conceal. The concept of separation of duties emerged, calling for the use of processing procedures that required more than one person to complete a transaction. As the books of account became increasingly important, accounting standards were developed, and they continue to evolve to this day. These standards served to make books of account comparable and to assure outsiders that an organization's books of account presented an accurate picture of its condition and assets. These developments led, in turn, to the requirement that an outside auditor perform an independent review of the books of account and operating procedures.

The transition to automated accounting systems introduced additional security requirements. Some early safeguards, such as the rule against erasures or changes in the books of account, no longer applied. Some computerized accounting systems lacked an audit trail, and others could have the audit trail subverted as easily as actual entries.

Finally, with the advent of the Information Age, intellectual property has become an increasingly important part of corporate and governmental assets. At the same time that intellectual property has grown in importance, threats to intellectual property have become more dangerous, because of information system (IS) technology itself. When sensitive information was stored on paper and other tangible documents, and rapid copying was limited to photography, protection was relatively straightforward. Nevertheless, document control systems, information classification procedures, and need-to-know access controls were not foolproof, and information compromises occurred with dismaying regularity. Evolution of IS technology has made information control several orders of magnitude more complex. The evolution and, more importantly, the implementation of control techniques have not kept pace.

The balance of this chapter describes how the evolution of information systems has caused a parallel evolution of information system security and at the same time has increased the importance of anticipating the impact of technical changes yet to come. This overview will clarify the factors leading to today's information system security risk environment and mitigation techniques and will serve as a warning to remain alert to the implication of technical innovations as they appear. The remaining chapters of this *Handbook* discuss information system security risks, threats, and vulnerabilities, their prevention and remediation, and many related topics in considerable detail.

1.2 EVOLUTION OF INFORMATION SYSTEMS.

The first electromechanical punched-card system for data processing, developed by Herman Hollerith at the end of the nineteenth century, was used to tabulate and total census field reports for the U.S. Bureau of the Census in 1890. The first digital, stored-program computers developed in the 1940s were used for military purposes, primarily cryptanalysis and the calculation and printing of artillery firing tables. At the same time, punched-card systems were already being used for accounting applications and were an obvious choice for data input to the new electronic computing machines.

1.2.1 1950s: Punched-Card Systems. In the 1950s, punched-card equipment dominated the commercial computer market.[2] These electromechanical devices could perform the full range of accounting and reporting functions. Because they were programmed by an intricate system of plugboards with a great many plug-in cables, and because care had to be exercised in handling and storing punched cards, only experienced persons were permitted near the equipment. Although any of these individuals could have set up the equipment for fraudulent use, or even engaged in sabotage, apparently few, if any, actually did so.

The punched-card accounting systems typically used four processing steps. As a preliminary, operators would be given a "batch" of documents, typically with an adding machine tape showing one or more "control totals." The operator keyed the data on each document into a punched card and then added an extra card, the batch control card, which stored the batch totals. Each card consisted of 80 columns, each containing, at most, one character. A complete record of an inventory item, for example, would be contained on a single card. The card was called a unit record, and the machines that processed the cards were called either unit record or punched-card machines. It was from the necessity to squeeze as much data as possible into an 80-character card that the later Year 2000 problem arose. Compressing the year into two characters was a universally used space-saving measure; its consequences 40 years later were not foreseen.

A group of punched cards, also called a "batch," were commonly held in a metal tray. Sometimes a batch would be rekeyed by a second operator, using a "verify-mode" rather than actually punching new holes in the cards, in order to detect keypunch errors before processing the card deck. Each batch of cards would be processed separately, so the processes were referred to as "batch jobs."

The first step would be to run the batch of cards through a simple program, which would calculate the control totals and compare them with the totals on the batch control card. If the batch totals did not reconcile, the batch was sent back to the keypunch area for rekeying. If the totals reconciled, the deck would be sort-merged with other batches of the same transaction type, for example, the current payroll. When this step was complete, the new batch consisted of a punched card for each employee in employee-number order. The payroll program accepted this input data card deck and processed the cards one by one. Each card was matched up with the corresponding employee's card in the payroll master deck to calculate the current net pay and itemized deductions and to punch a new payroll master card, including year-to-date totals. The final step was to use the card decks to print payroll checks and management reports. These steps were identical with those used by early, small-scale electronic computers. The only difference was in the speed at which the actual calculations were made. A complete process was still known as a batch job.

With this process, the potential for abuse was great. The machine operator could control every step of the operation. Although the data was punched into cards and verified by others, there was always a keypunch machine nearby for use by the machine operator. Theoretically, that person could punch a new payroll card and a new batch total card to match the change before printing checks and again afterward. The low incidence of reported exploits was due to the controls that discouraged such abuse, and possibly to the pride that machine operators experienced in their jobs.

1.2.2 Large-Scale Computers. While these electromechanical punched card machines were sold in large numbers, research laboratories and universities were working to design large-scale computers that would have a revolutionary effect on

the entire field. These computers, built around vacuum tubes, are known as the first generation. In March 1951, the first Universal Automatic Computer (UNIVAC) was accepted by the U.S. Census Bureau. Until then, every computer had been a one-off design, but UNIVAC was the first large-scale, mass-produced computer, with a total of 46 built. The word "universal" in its name indicated that UNIVAC was also the first computer designed for both scientific and business applications.[3]

UNIVAC contained 5,200 vacuum tubes, weighed 29,000 pounds, and consumed 125 kilowatts of electrical power. It dispensed with punched cards, receiving input from half-inch-wide metal tape recorded from keyboards, with output either to a similar tape or to a printer. Although not a model for future designs, its memory consisted of 1,000 72-bit words and was fabricated as a mercury delay line. Housed in a cabinet about six feet tall, two feet wide, and two feet deep was a mercury-filled coil running from top to bottom. A transducer at the top propagated slow-moving waves of energy down the coil to a receiving transducer at the bottom. There it was reconverted into electrical energy and passed on to the appropriate circuit, or recirculated if longer storage was required.

In 1956, IBM introduced the Random Access Method of Accounting and Control (RAMAC) magnetic disk system. It consisted of 50 magnetically coated metal disks, each 24 inches in diameter, mounted on a common spindle. Under servo control, two coupled read/write heads moved to span each side of the required disk and then inward to any one of 100 tracks. In one revolution of the disks, any or all of the information on those two tracks could be read out or recorded. The entire system was almost the size of a compact car and held what, for that time, was a tremendous amount of data—5 megabytes. The cost was $10,000 per megabyte, or $35,000 per year to lease. This compares with some of today's magnetic hard drives that measure about $3^1/_2$ inches wide by 1 inch high, store as much as 1,000 gigabytes, and cost less than $400, or about $0.0004 per megabyte.

Those early, massive computers were housed in large, climate-controlled rooms. Within the room, a few knowledgeable experts, looking highly professional in their white laboratory coats, attended to the operation and maintenance of their million-dollar charges. The concept of a "user" as someone outside the computer room who could interact directly with the actual machine did not exist.

Service interruptions, software errors, and hardware errors were usually not critical. If any of these caused a program to fail or abort, beginning again was a relatively simple matter. Consequently, the primary security concerns were physical protection of the scarce and expensive hardware, and measures to increase their reliability. Another issue, then as now, was human fallibility. Because the earliest computers were programmed in extremely difficult machine language, consisting solely of ones (1s) and zeros (0s), the incidence of human error was high, and the time to correct errors was excessively long. Only later were assembler and compiler languages developed to increase the number of people able to program the machines and to reduce the incidence of errors and the time to correct them.

Information system security for large-scale computers was not a significant issue then for two reasons. First, only a few programming experts were able to utilize and manipulate computers. Second, there were very few computers in use, each of which was extremely valuable, important to its owners, and consequently closely guarded.

1.2.3 Medium-Size Computers. In the 1950s, smaller computer systems were developed with a very simple configuration; punched-card master files were replaced by punched paper tape and, later, by magnetic tape, and disk storage systems. The

electromechanical calculator with its patchboard was replaced by a central processor unit (CPU) that had a small main memory, sometimes as little as 8 kilobytes,[4] and limited processing speed and power. One or two punched-card readers could read the data and instructions stored on that medium. Later, programs and data files were stored on magnetic tape. Output data were sent to cardpunches, for printing on unit record equipment, and later to magnetic tape. There was still no wired connection to the outside world, and there were no online users because no one, besides electronic data processing (EDP) people within the computer room, could interact directly with the system. These systems had very simple operating systems and did not use multiprocessing; they could run only one program at a time.

The IBM Model 650, as an example, introduced in 1954, measured about 5 feet by 3 feet by 6 feet and weighed almost 2,000 pounds. Its power supply was mounted in a similarly sized cabinet, weighing almost 3,000 pounds. It had 2,000 (10-digit) words of magnetic drum primary memory, with a total price of $500,000 or a rental fee of $3,200 per month. For an additional $1,500 per month, a much faster core memory, of 60 words, could be added. Input and output both utilized read/write punch-card machines. The typical 1950s IS hardware was installed in a separate room, often with a viewing window so that visitors could admire the computer. In an early attempt at security, visitors actually within the computer room were often greeted by a printed sign saying:

> Achtung! Alles Lookenspeepers!
>
> Das computermachine ist nicht fur gefingerpoken und mittengrabben.
> Ist easy schnappen der springenwerk, blowenfusen, und poppencorken mit spitzensparken.
> Ist nicht fur gewerken bei das dumbkopfen.
> Das rubbernecken sightseeren keepen hans in das pockets muss. Relaxen und watch das blinkenlichten.

Since there were still no online users, there were no user IDs and passwords. Programs processed batches of data, run at a regularly scheduled time—once a day, once a week, and so on, depending on the function. If the data for a program were not available at the scheduled run time, the operators might run some other job instead and wait for the missing data. As the printed output reports became available, they were delivered by hand to their end users. End users did not expect to get a continuous flow of data from the information processing system, and delays of even a day or more were not significant, except perhaps with paycheck production.

Information system security was hardly thought of as such. The focus was on batch controls for individual programs, physical access controls, and maintaining a proper environment for the reliable operation of the hardware.

1.2.4 1960s: Small-Scale Computers. During the 1960s, before the introduction of small-scale computers, dumb[5] terminals provided users with a keyboard to send a character stream to the computer and a video screen that could display characters transmitted to it by the computer. Initially, these terminals were used to help computer operators control and monitor the job stream, while replacing banks of switches and indicator lights on the control console. However, it was soon recognized that these terminals could replace card readers and keypunch machines as well. Now users, identified by user IDs, and authenticated with passwords, could enter input data through a CRT terminal into an edit program, which would validate the input and then store it on a hard drive until it was needed for processing. Later, it was realized that users also could directly access data stored in online master files.

1.2.5 Transistors and Core Memory. The IBM 1401, introduced in 1960 with a core memory of 4,096 characters, was the first all-transistor computer, marking the advent of the second generation. Housed in a cabinet measuring 5 feet by 3 feet, the 1401 required a similar cabinet to add an additional 12 kilobytes of main memory. Just one year later, the first integrated circuits were used in a computer, making possible all future advances in miniaturizing small-scale computers and in reducing the size of mainframes significantly.

1.2.6 Time Sharing. In 1961, the Compatible Time Sharing System (CTSS) was developed for the IBM 7090/7094. This operating system software, and its associated hardware, was the first to provide simultaneous remote access to a group of online users through multiprogramming.[6] "Multiprogramming" means that more than one program can appear to execute at the same time. A master control program, usually called an operating system (OS), managed execution of the functional applications programs. For example, under the command of the operator, the OS would load and start application #1. After 50 milliseconds, the OS would interrupt the execution of application #1 and store its current state in memory. Then the OS would start application #2 and allow it to run for 50 milliseconds, and so on. Usually, within a second after users had entered keyboard data, the OS would give their applications a time slice to process the input. During each time slice, the computer might execute hundreds of instructions. These techniques made it appear as if the computer were entirely dedicated to each user's program. This was true only so long as the number of simultaneous users was fairly small. After that, as the number grew, the response to each user slowed down.

1.2.7 Real-Time, Online Systems. Because of multiprogramming and the ability to store records online and accessible in random order, it became feasible to provide end users with direct access to data. For example, an airline reservation system stores a record of every seat on every flight for the next 12 months. A reservation clerk, working at a terminal, can answer a telephoned inquiry, search for an available seat on a particular flight, quote the fare, sell a ticket to the caller, and reserve the seat. Similarly, a bank officer can verify an account balance and effect money transfers. In both cases, each data record can be accessed and modified immediately, rather than having to wait for a batch to be run. Today, both the reservation clerk and the bank officer can be replaced by the customers themselves, who directly interface with the online computers.

While this advance led to a vast increase in available computing power, it also increased greatly the potential for breaches in computer security. With more complex operating systems, with many users online to sensitive programs, and with databases and other files available to them, protection had to be provided against inadvertent error and intentional abuse.

1.2.8 A Family of Computers. In 1964, IBM announced the S/360 family of computers, ranging from very small-scale to very large-scale models. All of the six models used integrated circuits, which marked the beginning of the third generation of computers. Where transistorized construction could permit up to 6,000 transistors per cubic foot, 30,000 integrated circuits could occupy the same volume. This lowered the costs substantially, and companies could buy into the family at a price within their means. Because all computers in the series used the same programming language and the same peripherals, companies could upgrade easily when necessary. The 360 family

quickly came to dominate the commercial and scientific markets. As these computers proliferated, so did the number of users, knowledgeable programmers, and technicians. Over the years, techniques and processes were developed to provide a high degree of security to these mainframe systems.

The year 1964 also saw the introduction of another computer with far-reaching influence: the Digital Equipment Corp. (DEC) PDP-8. The PDP-8 was the first mass-produced true minicomputer. Although its original application was in process control, the PDP-8 and its progeny quickly proved that commercial applications for minicomputers were virtually unlimited. Because these computers were not isolated in secure computer rooms but were distributed throughout many unguarded offices in widely dispersed locations, totally new risks arose, requiring innovative solutions.

1.2.9 1970s: Microprocessors. The foundations of all current personal computers (PCs) were laid in 1971 when Intel introduced the 4004 computer on a chip. Measuring 1/16 inch long by 1/8 inch high, the 4004 contained 2,250 transistors with a clock speed of 108 kiloHertz. The current generation of this earliest programmable microprocessor contains millions of transistors, with speeds over 1 gigaHertz, or more than 10,000 times faster. Introduction of microprocessor chips marked the fourth generation.

1.2.10 The First Personal Computers. Possibly the first personal computer was advertised in *Scientific American* in 1971. The KENBAK–1, priced at $750, had three programming registers, five addressing modes, and 256 bytes of memory. Although not many were sold, the KENBACK–1 did increase public awareness of the possibility for home computers.

It was the MITS Altair 8800 that became the first personal computer to sell in substantial quantities. Like the KENBACK–1, the Altair 8800 had only 256 bytes of memory, but it was priced at $375 without keyboard, display, or secondary memory. About one year later, the Apple II, designed by Steve Jobs and Steve Wozniak, was priced at $1,298, including a CRT display and a keyboard.

Because these first personal computers were entirely stand-alone and usually under the control of a single individual, there were few security problems. However, in 1978, the VisiCalc spreadsheet program was developed. The advantages of standardized, inexpensive, widely used application programs were unquestionable, but packaged programs, as opposed to custom designs, opened the way for abuse because so many people understood their user interfaces as well as their inner workings.

1.2.11 The First Network. A national network, conceived in late 1969, was born as ARPANET[7] (Advanced Research Projects Agency Network), a Department of Defense–sponsored effort to link a few of the country's important research universities, with two purposes: to develop experience in interconnecting computers and to increase productivity through resource sharing. This earliest connection of independent large-scale computer systems had just four nodes: the University of California at Los Angeles (UCLA), the University of California at Santa Barbara, Stanford Research Institute, and the University of Utah. Because of the inherent security in each leased-line interconnected node, and the physically protected mainframe computer rooms, there was no apparent concern for security issues. It was this simple network, with no thought of security designed in, from which evolved today's ubiquitous Internet and the World Wide Web (WWW), with their vast potential for security abuses.

1.2.12 Further Security Considerations. With the proliferation of remote terminals on commercial computers, physical control over access to the computer room was no longer sufficient. In response to the new vulnerabilities, logical access control systems were developed. An access control system maintains an online table of authorized users. A typical user record would store the user's name, telephone number, employee number, and information about the data the user was authorized to access and the programs the user was authorized to execute. A user might be allowed to view, add, modify, and delete data records in different combinations for different programs.

At the same time, system managers recognized the value of being able to recover from a disaster that destroyed hardware and data. Data centers began to make regular tape copies of online files and software for offsite storage. Data center managers also began to develop and implement offsite disaster recovery plans, often involving the use of commercial disaster-recovery facilities. Even with such a system in place, new vulnerabilities were recognized throughout the following years, and these are the subjects of much of this *Handbook.*

1.2.13 The First "Worm". A prophetic science-fiction novel, *The Shockwave Rider*, by John Brunner[8] (1975), depicted a "worm" that grew continuously throughout a computer network. The worm eventually exceeded a billion bits in length and became impossible to kill without destroying the network. Although actual worms (e.g., the Morris Worm of 1988) later became real-and-present menaces to all networked computers, prudent computer security personnel install constantly updated antimalware programs that effectively kill viruses and worms without having to kill the network.

1.2.14 1980s: Productivity Enhancements. The decade of the 1980s might well be termed the era of productivity enhancement. The installation of millions of personal computers in commercial, industrial, and government applications enhanced efficiency and functionality of vast numbers of users. These advances, which could have been achieved in no other way, were made at costs that virtually any business could afford.

1.2.15 1980s: The Personal Computer. In 1981, IBM introduced a general-purpose small computer it called the "Personal Computer." That model and similar systems became known generically as PCs. Until then, small computers were produced by relatively unknown sources, but IBM, with its worldwide reputation, brought PCs into the mainstream. The fact that IBM had demonstrated a belief in the viability of PCs made them serious contenders for corporate use.

There were many variations on the basic Model 5100 PC, and sales expanded far beyond IBM's estimates. The basic configuration used the Intel 8088, operating at 4.77 megaHertz, with up to two floppy disk drives, each of 160 kilobytes capacity and with a disk-based operating system (DOS) in an open architecture. This open OS architecture, with its available "hooks," made possible the growth of independent software producers, the most important of which was the Microsoft Corporation, formed by Bill Gates and Paul Allen.

IBM had arranged for Gates and Allen to create the DOS operating system. Under the agreement, IBM would not reimburse Gates and Allen for their development costs; rather, all profits from the sale of DOS would accrue to them. IBM did not have an exclusive right to the operating system, and Microsoft began selling it to many other customers as MS-DOS. IBM initially included with its computer the VisiCalc

spreadsheet program, but soon sales of Lotus 1-2-3 surpassed those of VisiCalc. The open architecture not only made it possible for many developers to produce software that would run on the PC, but also enabled anyone to put together purchased components into a computer that would compete with IBM's PC. The rapid growth of compatible application programs, coupled with the ready availability of compatible hardware, soon resulted in sales of more than 1 million units. Many subsequent generations of the original hardware and software are still producing sales measured in millions every year.

Apple took a very different approach with its Macintosh computer. Where IBM's system was wide open, Apple maintained tight control over any hardware or software designed to operate on the Macintosh so as to assure compatibility and ease of installation. The most important Apple innovations were the graphical user interface (GUI) and the mouse, both of which worked together to facilitate ease of use and both of which were derived from research and development at the Stanford Research Institute and the Xerox Palo Alto Research Institute in the 1960s and 1970s. Microsoft had attempted in 1985 to build these features into the Windows operating system, but early versions were generally rejected as slow, cumbersome, and unreliable. It was not until 1990 that Windows 3.0 overcame many of its problems and provided the foundation for later versions that were almost universally accepted.

1.2.16 Local Area Networks. During the 1980s, stand-alone desktop computers began to perform word processing, financial analysis, and graphic processing. Although this arrangement was much more convenient for end users than was a centralized facility, it was more difficult to share data with others.

As more powerful PCs were developed, it became practical to interconnect them so that their users could easily share data. These arrangements were commonly referred to as local area networks (LANs) because the hardware units were physically close, usually in the same building or office area. LANs have remained important to this day. Typically, a more powerful PC with a high storage capacity fixed[9] disk was designated as the file server. Other PCs, referred to as workstations, were connected to the file server using network interface cards installed in the workstations with cables between these cards and the file server. Special network software installed on the file server and workstations made it possible for workstations to access defined portions of the file server fixed disk just as if these portions were installed on the workstations. Furthermore, these shared files could be backed up at the file server without depending on individual users. By 1997, it was estimated that worldwide there were more than 150 million PCs operating as LAN workstations. The most common network operating systems (NOS) were Novell NetWare and later Microsoft IE (Internet Explorer).

Most LANs were implemented using the Ethernet (IEEE 802.3) protocol.[10] The server and workstations could be equipped with a modem (modulator/demodulator) connected to a dedicated telephone line. The modem enabled remote users, with a matching modem, to dial into the LAN and log on. This was a great convenience to LAN users who were traveling or working away from their offices, but such remote access created yet another new security issue. For the first time, computer systems were exposed in a major way to the outside world. From then on, it was possible to interact with a computer from virtually anywhere and from locations not under the same physical control as the computers themselves.

Typical NOS logical access control software provided for user IDs and passwords and selective authority to access file server data and program files. A workstation user logged on to the LAN by executing a log-in program resident on the file server.

The program prompted the user to enter an ID and password. If the log-in program concluded that the ID and password were valid, it consulted an access-control table to determine which data and programs the user might access. Access modes were defined as read-only, execute-only, create, modify (write or append), lock, and delete, with respect to individual files and groups of files. The LAN administrator maintained the access control table using a utility program. The effectiveness of the controls depended on the care taken by the administrator, and so, in some circumstances, controls could be weak. It was essential to protect the ID and password of the LAN administrator since, if they were compromised, the entire access-control system became vulnerable. Alert information system security officers noted that control over *physical* access to LAN servers was critical in maintaining the logical access controls. Intruders who could physically access a LAN server could easily restart the server using their own version of the NOS, completely bypassing the installed logical access controls.

Superficially, a LAN appears to be the same as a 1970s mainframe with remote dumb terminals. The difference technically is that each LAN workstation user is executing programs on the workstation, not on the centralized file server, while mainframe computers use special software and hardware to run many programs concurrently, one program for each terminal. To the user at a workstation or remote terminal, the two situations appear to be the same, but from a security standpoint, there are significant differences. The mainframe program software stays on the mainframe and cannot, under normal conditions, be altered during execution. A LAN program on a workstation can be altered, for example, by a computer virus, while actually executing. As a rule, mainframe remote terminals cannot download and save files, whereas workstations usually have at least a removable disk drive. Furthermore, a malicious workstation user can easily install a rewritable CD device, which makes it much easier to copy and take away large amounts of data.

Another important difference is the character of the connection between the computer and the terminals. Each dumb terminal has a dedicated connection to its mainframe and receives only that data that is directed to it. A LAN operates more like a set of radio transmitters sharing a common frequency on which the file server and the workstations take turns "broadcasting" messages. Each message includes a "header" block that identifies the intended recipient, but every node (the file server and the workstations) on a LAN receives all messages. Under normal circumstances, each node ignores messages not addressed to it. However, it is technically feasible for a workstation to run a modified version of the NOS that allows it to capture all messages. In this way, a workstation could identify all log-in messages and record the user IDs and passwords of all other users on the LAN, giving it complete access to all of the LAN's data and facilities.

Mainframe and LAN security also differ greatly in the operating environment. As noted, the typical mainframe is installed in a separate room and is managed by a staff of skilled technicians. The typical LAN file server, on the other hand, is installed in ordinary office space and is managed by a part-time, remotely located LAN administrator who may not be adequately trained. Consequently, the typical LAN has a higher exposure to tampering, sabotage, and theft. However, if the typical mainframe is disabled by an accident, fire, sabotage, or any other security incident, many business functions will be interrupted, whereas the loss of a LAN file server usually disrupts only a single function.

1.2.17 1990s: Interconnection. The Usenet evolved in the early 1980s as a free system for posting and retrieving news and commentary from participants—an

early form of disintermediation, since there were no controlling authorities to limit speech. Newsgroups developed on every conceivable topic, reaching tens of thousands of discussion areas within a few years. Computer enthusiasts and criminal hackers took to Usenet as an ideal channel for exchanging code, including details of hacks.

The commercial equivalents of the Usenet were the value-added networks (VANs) such as America On Line (AOL), CompuServe, and Prodigy. These services provided modems for telephone access, email, and facilities for defining discussion groups. Fees varied from hourly to monthly.

1.2.18 1990s: Total Interconnection. With the growing popularity of LANs, the technologies for interconnecting them emerged. These networks of physically interconnected local area networks were called wide area networks, or WANs. Any node on a LAN could access every node on any other interconnected LAN, and in some configurations, those nodes might also be given access to mainframe and minicomputer files and to processing capabilities.

1.2.19 Telecommuting. Once the WAN technology was in place, it became feasible to link LANs together by means of telecommunications circuits. It had been expensive to do this with the low-speed, online systems of the 1970s because all data had to be transmitted over the network. Now, since processing and most data used by a workstation were on its local LAN, a WAN network was much less expensive. Low-traffic LANs were linked using dial-up access for minimum costs, while major LANs were linked with high-speed dedicated circuits for better performance. Apart from dial-up access, all network traffic typically flowed over nonswitched private networks. Of the two methods, dial-up communications were considerably more vulnerable to security violations, and they remain so to this day.

1.2.20 Internet and the World Wide Web. The Internet, which began life in 1969 as the Advanced Research Projects Agency Network (ARPANET), slowly emerged onto the general computing scene during the 1980s. Initially, access to the Internet was restricted to U.S. Government agencies and their contractors. ARPANET users introduced the concept of email as a convenient way to communicate and exchange documents. Then, in 1989–1990, Sir Tim Berners-Lee conceived of the World Wide Web and the Web browser. This one concept produced a profound change in the Internet, greatly expanding its utility and creating an irresistible demand for access. During the 1990s, the U.S. Government relinquished its control, and the Internet became the gigantic, no-one-is-in-charge network of networks it is today. The explosive growth in participation in the global Internet is generally viewed as having started with the opening up of the .COM top-level domain to general use in 1993.

The Internet offers several important advantages: The cost is relatively low, connections are available locally in most industrialized countries, and by adopting the Internet protocol, TCP/IP, any computer becomes instantly compatible with all other Internet users.

The World Wide Web technology made it easy for anyone to access remote data. Almost overnight, the Internet became the key to global networking. Internet service providers (ISPs) operate Internet-compatible computers with both dial-up and dedicated access. A computer may access an ISP directly as a stand-alone ISP client or via a gateway from a LAN or WAN. A large ISP may offer dial-up access at many locations, sometimes called points of presence or POPs, interconnected by its own network. ISPs establish links with one another through the national access points (NAPs) initially set

up by the National Science Foundation. With this "backbone" in place, any node with access can communicate with another node, connected to a different ISP, located half way around the globe, without making prior arrangements.

The unrestricted access provided by the Internet created new opportunities for organizations to communicate with clients. A company can implement a Web server with a full-time connection to an ISP and open the Web server, and the WWW pages it hosts, to the public. A potential customer can access a Website, download product information and software updates, ask questions, and even order products. Commercial Websites, as they evolved from static "brochure-ware" to online shopping centers, stock brokerages, and travel agencies, to name just a few of the uses, became known as e-businesses.

1.2.21 Virtualization and the Cloud. As far back as the late 1960s, software was available to create encapsulated versions of an operating system on mainframe computers. Users interacted with what appeared to be their own, private mainframe environment. By the late 1980s, vendors created simulations of operating environments that could run under different operating systems (e.g., one could run DOS programs on UNIX machines). The trend continued throughout the succeeding years so that it is commonplace now to run programs under hypervisors that simulate complete or functionally limited versions of required operating systems on shared hardware.[11]

Today it is possible to provide users with instances of an operating environment on shared hardware, often at a distance, so that incremental increases in requirements can be satisfied at modest costs instead of having to purchase large-scale improvements in the hardware infrastructure. The situation is similar to what service bureaus offered in the decades when mainframe time-sharing was popular.

Another development in the last decade has been the availability of cloud computing, which refers to computer services, including storage (see Chapters 36 and 68), software as a service (SAAS), and infrastructure or platform as a service (IAAS and PAAS). See Chapter 68 for more details of managing and securing cloud computing.

1.2.22 Supervisory Control and Data Acquisition. The use of computers to control production of goods and services through supervisory control and data acquisition (SCADA) software and hardware has been growing throughout the four decades since this *Handbook* was first published in 1973. SCADA systems for critical infrastructure have been of great concern because contrary to initial design specifications, many of them have been connected to the general Internet, opening the systems they govern to subversion. For more about SCADA in information warfare, see Chapter 14 in this *Handbook.*

1.3 GOVERNMENT RECOGNITION OF INFORMATION ASSURANCE.

Certain major events in the history of information assurance (IA) center on government initiatives. In particular, IA has been strongly influenced by the development of security standards starting in the 1980s, by the publication of the landmark publication *Computers and Risk* in 1991, and by the establishment of the InfraGard program in the late 1990s for protection of the U.S. critical infrastructure.

1.3.1 IA Standards. In the late 1970s, the U.S. Department of Defense "established a Computer Security Initiative to foster the wide-spread availability of trusted

computer systems."[12] The author of the initial report that later became the *Trusted Computer Systems Evaluation Criteria* (TCSEC), DoD Standard 5200.28, wrote,

> Trusted computer systems are operating systems capable of preventing users from accessing more information than that to which they are authorized. Such systems are in great demand as more processing is entrusted to computers, while less information should be shared by all the system's users. With this demand comes a need to ascertain the integrity of computer systems on the market …

The TCSEC was issued with a bright orange cover and became known as the "Orange Book." Under the direction of National Computer Security Center (NCSC) director Patrick Gallagher and others, the National Security Agency (NSA) issued a "Rainbow Series" of books that profoundly affected the direction of IA in the USA and globally.[13]

The Rainbow Series led to similar efforts in other countries, culminating in the Common Criteria Evaluation and Validation Scheme (CCEVS), which has become the international standard for defining security levels for systems and software and for determining acceptable methods for testing and certifying system compliance with such standards.[14]

For details of the evolution of security standards, see Chapter 51 in this *Handbook*.

1.3.2 Computers at Risk.[15] In 1988, the Defense Advanced Research Projects Agency (DARPA) asked the Computer Science and Technology Board (renamed the Computer Science and Telecommunications Board of the NRC in 1990) for a study of computer and communications security issues affecting U.S. Government and industry. The NRC's System Security Study Committee published its results in a readable and informative book, *Computers at Risk: Safe Computing in the Information Age*.[16]

The Committee included experts with impeccable credentials, including executives from major computer vendors such as HP, DEC, and IBM; from high-technology companies such as Shearson, Lehman, Hutton Inc., and Rockwell International; universities such as Harvard and MIT; and think tanks like the RAND Corporation.

A public misconception is the supposed divergence in focus of the military and of commerce: The military is usually described as concerned with external threats and the problem of disclosure, whereas businesses are said to worry more about insider threats to data integrity. On the contrary, the military and commerce need to protect data in similar ways. The differences arise primarily from (1) the sophistication and resources available to governments that try to crack foreign military systems; (2) the relatively strong military emphasis on prevention compared with commercial need for proof that can be used in legal proceedings; and (3) the availability to the military of deep background checks on personnel, contrasted with the limits imposed on the invasion of privacy in the commercial sector.

Some of the more interesting points raised by the NRC Committee assert that:

- Because of the rapid and discontinuous pace of innovation in the computer field, "with respect to computer security, the past is not a good predictor of the future";
- Embedded systems (those where the microprocessor is not accessible to reprogramming by the user; e.g., medical imaging systems) open us to greater risks from inadequate quality assurance (e.g., a software bug in a Therac 25 linear accelerator killed three patients by irradiating them with more than 100 times the intended radiation dosage);

- Networking makes it possible to harm many more systems: "Interconnection gives an almost ecological flavor to security; it creates dependencies that can harm as well as benefit the community …"

The Committee proposed major recommendations, summarized as follows:

1. Push for implementation of generally accepted system security principles including:
   - Quality assurance standards that address security considerations;
   - Access control for operations as well as data (e.g., any of the menu systems which preclude access to the operating system);
   - Unambiguous user identification (ID) and authentication (e.g., personal profiles and hand-held password generators)
   - Protection of executable code (e.g., flags to show that certain object modules are "production" or "installed" and thus apply strict access control that would prevent unauthorized modification—as found in configuration control systems);
   - Security logging (e.g., logging failed file-open attempts and logon password violations);
   - Assigning a security administrator to each enterprise;
   - Data encryption;
   - Operational support tools for verifying the state and effectiveness of security measures (e.g., audit tools);
   - Independent audits of system security by people not directly involved in programming or system management of the audited system;
   - Hazard analysis evaluating threats to safety from different malfunctions and breaches of security (e.g., consequences of tampering with patient data in hospitals).
2. Take specific short-term actions now:
   - Develop security policies for your organization before there's a problem;
   - Form and train computer emergency response teams before a crisis to respond to security violations or attacks;
   - Use the Orange Book's (TCSEC, from the National Computer Security Center's Rainbow series) C2 and B1 criteria to define guidelines on security;
   - Improve software systems development by applying better quality-assurance methods;
   - Contribute to voluntary industry groups developing modern security standards and implement those standards in commercial software;
   - Make effective security the default in software and hardware (make the user explicitly disable security instead of having to enable it).
3. Learn and teach about security:
   - Build a repository of incident data;
   - Foster education in engineering secure systems, both by encouraging universities to provide postgraduate training in security and by urging industry to include security training as part of software engineering projects;

- Teach beginners about security and ethics in computer usage and programming (e.g., the NCSA is working on a research and development project to study beliefs, attitudes, and behavior about ethical issues in computing in grade and high schools, colleges, and universities).

4. Clarify export control criteria and set up a forum for arbitration (hardware and software vendors have been complaining for years that the arbitrary imposition of severe export restrictions hampers American competitiveness in overseas markets without materially helping national security).
5. Fund and pursue needed research in such areas as:
   - Security modularity: the effects on security of combining modules with known security properties;
   - Security policy models: more subtle requirements like integrity and availability are still not easily represented by control structures;
   - Cost estimation: there should be better ways of measuring the costs and benefits of security mechanisms in particular applications;
   - New technology: networking, in particular, leads to greater complexity (e.g., how to connect "mutually suspicious organizations");
   - Quality assurance for security: how to measure effectiveness;
   - Modeling tools: standards for graphical representations of security relationships analogous to the diagrams used in functional decomposition and object-oriented methodologies for program design;
   - Automated procedures: audit and monitoring tools for the data center management team;
   - Nonrepudiation: combining the need for detailed records of user actions with the values of privacy;
   - Resource control: how to ensure that proprietary software and data are used legitimately (e.g., preventing more than the licensed number of users from accessing a system, preventing software theft);
   - Security perimeters: how to reconcile the desire for network interconnection with limitations due to security requirements ("If, for example, a network permits mail but not directory services … less mail may be sent because no capability exists to look up the address of a recipient").

Chapter 2 of the NRC report, Concepts of Information Security, is a 25-page primer on information systems security that could be handed to any manager who needs to be filled in on why you propose to spend so much money protecting the computer systems. The authors cover the fundamental aspects of information security (confidentiality, integrity, and availability); management controls (individual accountability, auditing, and separation of duties); risks (probabilities of attack or damage) and vulnerabilities (weak points); and privacy issues. In Appendix 2.2, the authors report an informal survey in April 1989 of 30 private companies in a variety of fields. The consensus among those polled included the following basic standards for information systems security (show these to your upper management if necessary):

- Unique IDs, block access after a maximum number of incorrect logon attempts, show last successful access at logon time, make passwords and IDs expire;

- Disallow embedded passwords during logon, make passwords invisible during entry, force minimum length (6), store passwords encrypted, scan proposed passwords to eliminate easy words;
- Permit strict control over file access;
- Detect and interdict viruses, certify software as virus-free, provide data encryption, overwrite deleted files to prevent recovery, force tight binding of production data to production programs;
- Automated time-out for inactive sessions, unique identification of terminals and workstations during logon;
- Network security monitoring, modem locking, callback, automatic data encryption during transmission;
- Audit trails, including security violations;
- Generally applicable security standards that could be used by vendors and users to evaluate different equipment and software for specific environments.

Twenty years later, focus among information assurance experts has shifted beyond the technical to emphasize organizational controls. For example, the 2003 survey of members of the Information Systems Security Association included these information security function practices among the respondents:

- Access controls: 73%
- Written information security policy: 72%
- Compliance with existing laws and regulations: 66%
- Creation of organization and process to implement policy: 59%
- Awareness and training program: 57%
- Regular monitoring, reviewing, and auditing: 57%
- Business continuity planning: 57%
- Risk assessment and risk management: 56%

In 2007, Gary S. Miliefsky proposed the following seven priorities for corporate information security:

1. Policies
2. Awareness and training
3. Information security self-assessments
4. Regulatory compliance self-assessments
5. Corporate-wide encryption
6. Manage all corporate assets
7. Test Business Continuity Planning (BCP) and Disaster Recovery Planning (DRP)[17]

The Computer Security Division of the Information Technology Laboratory at the National Institute of Standards and Technology issued a draft reference model that

included the following "programmatic, integration, and system security activities that are typically a part of an information security program":

- Program Security Activities
  - Annual and Quarterly Review and Reporting of Information Security Program
  - Asset Inventory
  - Awareness and Specialized Security Training
  - Continuity of Operations
  - Incident Response
  - Periodic Testing and Evaluation
  - Plan of Action and Milestones
  - Policies and Procedures
  - Risk Management
- Integration Activities
  - Business Risk
  - Capital Planning and Investment Control (CPIC)
  - Configuration Management
  - Enterprise Architecture (EA)
  - Environmental Protection
  - Human Resources
  - Personnel Security
  - Physical Security
  - Privacy
  - Records Management
  - Strategic Plan
  - System Development Life Cycle (SDLC)
- System Security Activities
  - Categorize the Information System
  - Select Security Controls
  - Supplement Security Controls
  - Document Security Controls
  - Implement Security Controls
  - Assess Security Controls
  - Authorize the Information System
  - Monitor Security Controls

1.3.3 InfraGard.[18] InfraGard is a nationwide program in the United States that brings together representatives from information technology departments in industry and academia for information sharing and analysis, especially to help protect critical infrastructure against cyberattacks and also to support the FBI in its cybercrime investigations and education projects.[19]

The organization started in the Cleveland Field Office of the FBI in 1996 and expanded rapidly until there are now over 11,000 members in over 40 chapters. Joining InfraGard is easy and free for U.S. citizens residing in the United States. Using the Website (www.infragard.org), you can locate a nearby local chapter ("Find Chapters") and contact your chapter officers. You can get application forms online and then send them in to the FBI liaison officer for that chapter to be vetted for admission. The FBI conducts a background check to ensure that all members are likely to be trustworthy to participate in confidential discussions of threats and vulnerabilities. Chapters usually conduct regular local meetings and organize list-servers for exchange of information among members. Many have newsletters as well.

1.4 RECENT DEVELOPMENTS. In recent years, a key development has been the dramatic increase in availability of inexpensive portable data storage devices. At the time of writing (2013), flash drives the size of a lipstick or even of an antacid pill are available with capacities in the dozens of gigabytes for a few dollars. Such devices are available in a wide range of concealable formats such as pens, music players, watches, and (no joke) sushi. Pocket-sized hard disks and solid-state drives with capacities in the hundreds of gigabytes to terabytes are available for less than US$100. Digital cameras use storage cards that can be used for data transfers; mobile phones include cameras and recording capabilities. A 64 GB micro SD card for a phone costs about $50 and can hold 6,000 songs from iTunes—or the entire customer database being stolen by a disaffected soon-to-be-fired employee. Controlling data leakage through unauthorized connection of such devices has become a significant problem for security managers. Systems for restricting connection of devices and controlling data transfers to such storage media (data-loss prevention or DLP) are spreading through government and corporate environments (see Chapter 13 in this *Handbook* for a detailed discussion of DLP).

Another issue that increasingly concerns security managers is the protection of personally identifiable information (PII) from customers or data subjects. Many organizations including government agencies, banks, and universities have suffered serious damage from loss of control over PII and the risks of identity theft resulting from exposure of such sensitive data. Legislators are responding to public concern by increasing legal requirements for protection of PII. The use of encryption on mobile data systems such as laptop computers, personal digital assistants (PDAs), mobile phones, and integrated systems that combine many functions (e.g., BlackBerries) has become a necessity. See Chapter 69 in this *Handbook* for extensive discussion of protection of PII.

A consequence of the growing interconnectivity of storage and communications devices is that corporate networks are no longer insulated from less-secure systems. Users who connect poorly protected laptops (or other devices) to public networks such as hotel-supplied ISPs or wireless access points in coffee shops may return to their home offices with malware-infected systems that contaminate the entire network. Security managers are increasingly turning to integrated systems for controlling connectivity via virtual private networks and supervisory software that monitors and restricts unauthorized connections, software installations, and downloads.

Another growing issue is the increasing speed and persistence of attacks on systems and networks by state-sponsored and criminal organizations engaged in industrial espionage and fraud. See Chapters 2, 14, 15, and 46, among many others in this *Handbook*, for further discussion of the changing threat profile for today's information systems.

1.5 ONGOING MISSION FOR INFORMATION SYSTEM SECURITY.

There is no end in sight to the continuing proliferation of Internet nodes, to the variety of applications, to the number and value of online transactions, and, in fact, to the rapid integration of computers into virtually every facet of our existence. Nor will there be any restrictions as to time or place. With 24/7/365, always-on operation, and with global expansion even to relatively undeveloped lands, both the beneficial effects and the security violations can be expected to grow apace.

Convergence, which implies computers, televisions, cell phones, and other means of communication combined in one unit, together with continued growth of information technology, will lead to unexpected security risks. Distributed denial-of-service (DDoS) attacks, copyright infringement, child pornography, fraud, theft of identity, and industrial espionage are all ongoing security threats. So far, no perfect defensive measures have been developed.

The situation is currently (2013) changing from identifying vulnerabilities and preventing penetrations to identifying compromises and minimizing damage from long-lasting subversion of protection mechanisms. Situational awareness and rapid response are becoming an increasingly important element in long-term defenses of our information.

This *Handbook* provides a foundation for understanding and blunting both the existing vulnerabilities and those new threats that will inevitably arise in the future. Certainly, no one but the perpetrators could have foreseen the use of human-guided missiles to attack the World Trade Center. Besides its symbolic significance, the great concentration of resources within the WTC increased its attractiveness as a target. After 9/11, the importance of physical safety of personnel has become the dominant security issue, with disaster recovery of secondary, but still great, concern. This *Handbook* cannot foresee all possible future emergencies, but it does prescribe some preventative measures, and it does recommend procedures and resources for mitigation and remediation.

1.6 NOTES

1. Many technical specialists tend to use the term “security” to refer to logical access controls. A glance at the contents pages of this volume shows the much broader scope of information system security.
2. For further details, see, for example, www.cs.uiowa.edu/~jones/cards.
3. See http://ei.cs.uiowa.edu/~history/UNIVAC.Weston.html and inventors.about.com/library/weekly/aa062398.htm
4. It is notable that the IBM 1401 computer was so named because the initial model had 1,400 bytes of main memory. It was not long before memory size was raised to 8 kilobytes and then later to as much as 32 kilobytes. In 1980, the Series III minicomputer from Hewlett-Packard doubled its maximum memory from 1 megabyte to 2 megabytes at a cost of $64,000 (about $200,000 in 2008 dollars). This compares with today’s personal computers, typically equipped with no less than 512 megabytes and often a gigabyte or more.
5. The term “dumb” was used because the terminal had no internal storage or processing capability. It could only receive and display characters and accept and transmit keystrokes. Both the received characters and the transmitted ones were displayed on a cathode ray tube (CRT) much like a pre-color television screen. Consequently, these were also called “glass” terminals.

6. "Multiprocessing," "multiprogramming," and "multitasking" are terms that are used almost interchangeably today. Originally, multitasking implied that several modules or subroutines of a single program could execute together. Multiprogramming was designed to execute several different programs, and their subroutines, concurrently. Multiprocessing most often meant that two or more computers worked together to speed program execution by providing more resources.
7. Also known as ARPAnet and Arpanet.
8. First published in 1975; reissued by Mass Market Paperbacks in May 1990.
9. "Fixed," in contrast with the removable disk packs common in large data centers.
10. See standards.ieee.org/getieee802/802.3.html
11. Sean P. Conroy, "History of Virtualization," *Everything VM,* 2010. www.everythingvm.com/content/history-virtualization
12. G. H. Nibaldi, Proposed Technical Evaluation Criteria for Trusted Computer Systems. Publication M79-225 (Bedford, MA: MITRE Corporation, 1999).
13. For access to all the Rainbow Series documents, see www.fas.org/irp/nsa/rainbow.htm
14. The CCEVS Website has extensive documentation; see www.niap-ccevs.org
15. This section is reprinted with slight modifications by permission of the author from the original manuscript for M. E. Kabay, *The NCSA Guide to Enterprise Security: Protecting Information Assets* (New York: McGraw-Hill, 1996), Chapter 1, pp. 2–5.
16. National Research Council, *Computers at Risk: Safe Computing in the Information Age* (Washington, DC: National Academy Press, 1991). Available as searchable openbook at www.nap.edu/openbook.php?isbn=0309043883
17. G. S. Miliefsky, "The 7 Best Practices for Network Security in 2007," NetworkWorld Website, 2007, www.networkworld.com/columnists/2007/011707miliefsky.html?t51hb
18. M. E. Kabay (2005). "InfraGard is not a Deodorant," NetworkWorld Website, 2005, www.networkworld.com/newsletters/sec/2005/0905sec2.html
19. www.infragard.org

CHAPTER 2

HISTORY OF COMPUTER CRIME

M.E. Kabay

2.1 WHY STUDY HISTORICAL RECORDS? 2·2

2.2 OVERVIEW 2·2

2.3 1960S AND 1970S: SABOTAGE 2·3

2.3.1 Direct Damage to Computer Centers 2·3
2.3.2 1970–1972: Albert the Saboteur 2·4

2.4 IMPERSONATION 2·5

2.4.1 1970: Jerry Neal Schneider 2·5
2.4.2 1980–2003: Kevin Mitnick 2·5
2.4.3 Credit Card Fraud 2·6
2.4.4 Identity Theft Rises 2·7

2.5 PHONE PHREAKING 2·8

2.5.1 2600 Hz 2·8
2.5.2 1982–1991: Kevin Poulsen 2·8

2.6 DATA DIDDLING 2·9

2.6.1 Equity Funding Fraud (1964–1973) 2·9
2.6.2 1994: Vladimir Levin and the Citibank Heist 2·10

2.7 SALAMI FRAUD 2·10

2.8 LOGIC BOMBS 2·11

2.9 EXTORTION 2·12

2.10 TROJAN HORSES 2·12

2.10.1 1988 Flu-Shot Hoax 2·12
2.10.2 Scrambler, 12-Tricks, and PC Cyborg 2·12
2.10.3 1994: Datacomp Hardware Trojan 2·13
2.10.4 Keylogger Trojans 2·13
2.10.5 Haephrati Trojan 2·14
2.10.6 Hardware Trojans and Information Warfare 2·14

2.11 NOTORIOUS WORMS AND VIRUSES 2·15

2.11.1 1970–1990: Early Malware Outbreaks 2·15
2.11.2 December 1987: Christmas Tree Worm 2·16
2.11.3 November 2, 1988: Morris Worm 2·16
2.11.4 Malware in the 1990s 2·17
2.11.5 March 1999: Melissa 2·18
2.11.6 May 2000: I Love You 2·20
2.11.7 July 2010 Stuxnet 2·20

2.12 SPAM 2·20

2.12.1 1994: Green Card Lottery Spam 2·20
2.12.2 Spam Goes Global 2·21

2.13 DENIAL OF SERVICE 2·21

2.13.1 1996: Unamailer 2·21
2.13.2 2000: MafiaBoy 2·22

2.14 HACKER UNDERGROUND 2·22

2.14.1 1981: Chaos Computer Club 2·23

2.14.2 1982: The 414s 2·23
2.14.3 1984: Cult of the Dead Cow 2·23
2.14.4 1984: *2600: The Hacker Quarterly* 2·24
2.14.5 1984: Legion of Doom 2·24
2.14.6 1985: *Phrack* 2·25
2.14.7 1989: Masters of Deception 2·25
2.14.8 1990: Operation Sundevil 2·26
2.14.9 1990: Steve Jackson Games 2·26
2.14.10 1992: L0pht Heavy Industries 2·27
2.14.11 2004: Shadowcrew 2·27
2.14.12 Late 2000s: Russian Business Network (RBN) 2·27
2.14.13 Anonymous 2·28
2.14.14 2013: Unlimited Operations 2·28

2.15 INDUSTRIAL ESPIONAGE 2·29

2.16 CONCLUDING REMARKS 2·31

2.17 FURTHER READING 2·32

2.18 NOTES 2·33

2.1 WHY STUDY HISTORICAL RECORDS? Every field of study and expertise develops a common body of knowledge that distinguishes professionals from amateurs. One element of that body of knowledge is a shared history of significant events that have shaped the development of the field. Newcomers to the field benefit from learning the names and significant events associated with their field so that they can understand references from more senior people in the profession, and so that they can put new events and patterns into perspective. This chapter provides a brief overview of some of the more famous (or notorious) cases of computer crime (including those targeting computers and those mediated through computers) of the last four decades.[1]

2.2 OVERVIEW. This chapter illustrates several general trends from the 1960s through mid-2013:

- In the early decades of modern information technology (IT), computer crimes were largely committed by individual disgruntled and dishonest employees.
- Physical damage to computer systems was a prominent threat until the 1980s.
- Criminals often used authorized access to subvert security systems as they modified data for financial gain or destroyed data for revenge.
- Early attacks on telecommunications systems in the 1960s led to subversion of the long-distance phone systems for amusement and for theft of services.
- As telecommunications technology spread throughout the IT world, hobbyists with criminal tendencies learned to penetrate systems and networks.
- Programmers in the 1980s began writing malicious software, including self-replicating programs, to interfere with personal computers.
- As the Internet increased access to increasing numbers of systems worldwide, criminals used unauthorized access to poorly protected systems for vandalism, political action, and financial gain.
- As the 1990s progressed, financial crime using penetration and subversion of computer systems increased.
- The types of malware shifted during the 1990s, taking advantage of new vulnerabilities and dying out as operating systems were strengthened, only to succumb to new attack vectors.

- Illegitimate applications of email grew rapidly from the mid-1990s onward, generating torrents of unsolicited commercial and fraudulent email.
- Organized crime became increasingly involved in systematic penetration of financial systems and targeted fraud.
- Chinese government-supported civilian and military agents increasingly used computer-based industrial espionage to gain significant economic advantages over industry and commerce in North America and Europe.

2.3 1960s AND 1970s: SABOTAGE. Early computer crimes often involved physical damage to computer systems and subversion of the long-distance telephone networks.

2.3.1 Direct Damage to Computer Centers. In February 1969, the largest student riot in Canada was set off when police were called in to put an end to a student occupation of several floors of the Hall Building. The students had been protesting against a professor accused of racism, and when the police came in, a fire broke out and computer data and university property were destroyed. The damages totaled $2 million, and 97 people were arrested.[2]

Thomas Whiteside cataloged a litany of early physical attacks on computer systems in the 1960s and 1970s[3]:

1968 Olympia, WA: An IBM 1401 in the state is shot twice by a pistol-toting intruder
1970 University of Wisconsin: Bomb kills one and injures three people and destroys $16 million of computer data stored on site
1970 Fresno State College: Molotov cocktail causes $1 million damage to computer system
1970 New York University: Radical students place fire-bombs on top of Atomic Energy Commission computer in attempt to free a jailed Black Panther
1972 Johannesburg, South Africa: Municipal computer is dented by four bullets fired through a window
1972 New York: Magnetic core in Honeywell computer attacked by someone with a sharp instrument, causing $589,000 of damage
1973 Melbourne, Australia: Antiwar protesters shoot American firm's computer with double-barreled shotgun
1974 Charlotte, NC: Charlotte Liberty Mutual Life Insurance Company computer is shot by a frustrated operator
1974 Dayton, OH: Wright Patterson Air Force Base: Four attempts are made to sabotage computers, including by magnets, loosened wires, and gouges in equipment
1977 Rome, Italy: Four terrorists pour gasoline on university computer and burn it to cinders
1978 Lompoc, CA: Vandenburg Air Force Base: A peace activist destroys an unused IBM 3031 using a hammer, a crowbar, a bolt cutter, and a cordless power drill as a protest against the NAVSTAR satellite navigation system, claiming it gives the United States a first-strike capability

The incidents of physical abuse of computer systems did not stop as other forms of computer crime increased. For example, in 2001, *NewsScan* editors[4] summarized a report from *Wired Magazine*:

> A survey by British PC maker Novatech, intended to take a lighthearted look at techno-glitches, instead revealed the darker side of computing. One in every four computers has been physically assaulted by its owner, according to the 4,200 respondents.[5]

In April 2003, the National Information Protection Center and Department of Homeland Security reported:

> Nothing brings a network to a halt more easily and quickly than physical damage. Yet as data transmission becomes the lifeblood of Corporate America, most big companies haven't performed due diligence to determine how damage-proof their data lifelines really are. Only 20 percent of midsize and large companies have seriously sussed out what happens to their data connections after they go beyond the company firewall, says Peter Salus of MatrixNetSystems, a network-optimization company based in Austin, TX.[6]

By the mid-2000s, concerns over the physical security of electronic voting systems had risen to public awareness. For example:

> A cart of Diebold electronic voting machines was delivered today to the common room of this Berkeley, CA, boarding house, which will be a polling place on Tuesday's primary election. The machines are on a cart which is wrapped in plastic wrap (the same as the stuff we use in the kitchen). A few cable locks (bicycle locks, it seems) provide the appearance of physical security, but they aren't threaded through each machine. Moreover, someone fiddling with the cable locks, I am told, announced after less than a minute of fiddling that he had found the three-digit combination to be the same small integer repeated three times.[7]

2.3.2 1970–1972: Albert the Saboteur. One of the most instructive early cases of computer sabotage occurred at the National Farmers Union Service Corporation of Denver, where a Burroughs B3500 computer suffered 56 disk head crashes in the two years from 1970 to 1972. Downtime was as long as 24 hours per crash, with an average of 8 hours per incident. Burroughs experts were flown in from all over the United States at one time or another, and concluded that the crashes must be due to power fluctuations.

By the time all the equipment had been repaired and new wiring, motor generators, circuit breakers, and power-line monitors had been installed in the computer room, total expenditures for hardware and construction were over $500,000 (in 1970 dollars). Total expenses related to down time and lost business opportunities because of delays in providing management with timely information are not included in this figure. In any case, after all this expense, the crashes continued sporadically as before.

By this time, the experts were beginning to wonder about their analysis. For one thing, all the crashes had occurred at night. Could it be sabotage? Surely not! Old Albert, the night-shift operator, had been so helpful over all these years; he had unfailingly called in the crashes at once, gone out for coffee and donuts for the repair crews, and been meticulous in noting the exact times and conditions of each crash. However, all the crashes had in fact occurred on his shift.

Management installed a closed-circuit television (CCTV) camera in the computer room—without informing Albert. For some days, nothing happened. Then one night another crash occurred. On the CCTV monitor, security guards saw good ol' Albert open up a disk cabinet and poke his car key into the read/write head solenoid, shorting it out and causing the 57th head crash.

The next morning, management confronted Albert with the film of his actions and asked for an explanation. Albert broke down in mingled shame and relief. He confessed to an overpowering urge to shut the computer down. Psychological investigation determined that Albert, who had been allowed to work night shifts for years without a change, had simply become lonely. He arrived just as everyone else was leaving; he left as everyone else was arriving. Hours and days would go by without the slightest human

interaction. He never took courses, never participated in committees, never felt involved with others in his company. When the first head crashes occurred—spontaneously—he had been surprised and excited by the arrival of the repair crew. He had felt useful, bustling about, telling them what had happened. When the crashes had become less frequent, he had involuntarily, and almost unconsciously, re-created the friendly atmosphere of a crisis team. He had destroyed disk drives because he needed company.[8]

2.4 IMPERSONATION. Using the insignia and specialized language of officials as part of social engineering has a long history in crime; a dramatization of these techniques is in the popular movie *Catch Me If You Can*[9] about Frank William Abagnale Jr., the teenage scammer and counterfeiter who pretended to be a pilot, a doctor, and a prosecutor before eventually becoming a major contributor to the U.S. Government's anticounterfeiting efforts and then founding a major security firm.[10]

Several criminals involved in computer-mediated or computer-oriented crime became notorious for using impersonation.

2.4.1 1970: Jerry Neal Schneider. A notorious computer-related crime started in 1970, when teenager Jerry Neal Schneider used Dumpster® diving to retrieve printouts from the Pacific Telephone and Telegraph (PT&T) company in Los Angeles. After years of collection, he had enough knowledge of procedures that he was able to impersonate company personnel on the phone. He collected yet more detailed information on procedures. Posing as a freelance magazine writer, he even got a tour of the computerized warehouse and information about ordering procedures. In June 1971, he ordered $30,000 of equipment to be sent to a normal PT&T dropoff point—and promptly stole it and sold it. He eventually had a 6,000-square-foot warehouse and 10 employees. He stole over $1 million of equipment—and sold some of it back to PT&T. He was finally denounced by one of his own disgruntled employees and became a computer security consultant after his prison term.[11]

2.4.2 1980–2003: Kevin Mitnick. Born in 1963, Kevin Mitnick became involved in crime early, using a special punch for bus transfers to get free rides anywhere in the San Fernando Valley in California by the time he was a young teenager. His own autobiographical comments show him to have been involved in phone phreaking, malicious pranks, and breaking into computers at the Digital Equipment Corporation (DEC) using social engineering.[12]

In 1981, he and his friend Lewis De Payne used social engineering to gain unauthorized access to an operations center for Pacific Bell; "the juvenile court ordered a diagnostic psychological study of Mitnick and sentenced him to a year's probation."[13] In 1987, he was arrested for breaking into the computers of the Santa Cruz Operation, makers of SCO UNIX, and sentenced to three years' probation.

In the summer of 1988, Mitnick and his accomplice and friend Lenny DiCicco cracked the University of Southern California computers again and misappropriated hundreds of Mb of disk space (a lot at the time) to store VAX VMS source files stolen from Digital Equipment Corporation (DEC). Mitnick was arrested by the Federal Bureau of Investigation (FBI) for having stolen the VAX VMS source code. During his trial, he was described as suffering from an impulse-control disorder. In July 1989, he was sentenced to a year in jail and six months' rehabilitation. He later tried to become a private investigator and security specialist. He was generally treated with hostility by the established information security community.

In November 1992, Mitnick went underground again when the FBI got a warrant for his arrest on charges of stealing computer time from a phone company. He was located two years later when he made the mistake of leaving insulting messages on the computer and voicemail systems of a physicist and Internet security expert Tsutomu Shimomura. Shimomura was so irritated that he helped law enforcement authorities track the fugitive to North Carolina, where Mitnick was arrested in February 1995 and imprisoned pending trial.

Mitnick was convicted in federal court for the Central District of California on August 9, 1999, and sentenced to 46 months imprisonment for "four counts of wire fraud, two counts of computer fraud, and one count of illegally intercepting a wire communication."[14] Mitnick was previously sentenced by Judge Pfaelzer to an additional 22 months in prison, this for possessing cloned cellular phones when he was arrested in North Carolina in 1995, and for violating terms of his supervised release imposed after being convicted of an unrelated computer fraud in 1989. He admitted to violating the terms of supervised release by hacking into PacBell voicemail and other systems, and to associating with known computer hackers, in this case codefendant Louis De Payne. Following his release from prison in September 2000, Mitnick was to be on three years' parole, during which his access to computers was restricted[15] and his profits from writing or speaking about his criminal career were to be turned over to reimburse his victims.

Mitnick earned a living on the talk circuit and eventually founded his own security consulting firm. In the years since his release from prison, he has collaborated in writing several books on social engineering.[16]

Perhaps his most significant position in the history of computer crime is that he became an icon in the criminal underground. "FREE KEVIN" was a popular component of Web vandalism for many years, and Eric Corley, the longtime editor of the criminal-hacking publication *2600: The Hacker Quarterly,* even made a movie, *Freedom Downtime,* about what the criminal underground describes as the grossly unfair treatment of Mitnick by the federal government and the news media.[17]

2.4.3 Credit Card Fraud. Credit at local businesses dates back into the undocumented past.[18] In the United States, credit cards appeared in the mid-1920s when gasoline companies began issuing cards that were recognized at stations across the country.[19] In 1950, Frank X. McNamara started the Diners Club, the first credit card company serving multiple types of businesses; the company began the practice of charging a percentage fee for each transaction and also charged its clients a membership fee.[20] The VISA card evolved from the 1951 BankAmericard from the Bank of America, and a consortium of California banks established MasterCard shortly thereafter. American Express started its card program in 1958.

Card use rose and, unsurprisingly, credit card fraud was rampant. Mail theft also became widespread as unscrupulous individuals discovered that envelopes containing credit cards were just like envelopes full of cash. And there was little to stop card companies from sending out cards that customers had never asked for, were not expecting, and could not have known had been stolen, until the issuing company began demanding payment for the charges that had been run up. These crimes and other problems stemming from the relentless card-pushing by banks led directly to the passage of the Fair Credit Billing Act of 1974[21] as well as many other laws[22] designed to protect the consumer.[23]

By the mid-1990s, credit card fraud was a rapidly growing problem for consumers and for law enforcement. A 1997 FBI report stated:

> Around the world, bank card fraud losses to Visa and Master-Card alone have increased from $110 million in 1980 to an estimated $1.63 billion in 1995 … . The United States has suffered the bulk of these losses—approximately $875 million for 1995 alone. This is not surprising because 71 percent of all worldwide revolving credit cards in circulation were issued in this country … . Law enforcement authorities continually confront new and complex schemes involving credit card frauds committed against financial institutions and bank card companies. Perpetrators run the gamut from individuals with easy access to credit card information—such as credit agency officials, airline baggage handlers, and mail carriers, both public and private, to organized groups, usually from similar ethnic backgrounds, involved in large-scale card theft, manipulation, and counterfeiting activities. Although current bank card fraud operations are numerous and varied, several schemes account for the majority of the industry's losses by taking advantage of dated technology, customer negligence, and laws peculiar to the industry.[24]

2.4.4 Identity Theft Rises. By the late 1990s and in the decade following the year 2000, credit card fraud was subsumed into the broader category of *identity theft*. Instead of limiting their depredations to running up bills on stolen or forged credit card accounts, thieves, often in organized rings, created entire bogus parallel identities, initiating unpaid bank loans, buying cars with other people's credit, and wreaking havoc with innocent victims' credit ratings, financial situations, and even their daily life. Victims of extreme cases lost their ability to obtain mortgages, buy new homes, and accept new jobs. Worse, the burden of proof of innocence fell on the victims, in a bitter reversal of the assumption of innocence underlying British common law and its offshoot in the commonwealth and the United States.

In August 2008, the U.S. Department of Justice announced[25] the single largest and most complex case of identity theft ever charged in this country. It involved eleven people from five different countries, including two from the United States and two from the People's Republic of China, who had stolen more than 40,000,000 credit card records from a major U.S. retailer. They drove by, or loitered at, buildings in which wireless networks were housed, and installed sniffers that recorded passwords, card numbers, and account data. Unless adequate preventative measures are installed quickly, more such horrendous events will be sure to occur. For more on wireless network security, see Chapter 33 in this *Handbook*.

The 2011 report from the U.S. Bureau of Justice Statistics, *Identity Theft Reported by Households, 2005–2010* provides additional details. The abstract includes these highlights:

- In 2010, 7.0 percent of households in the United States, or about 8.6 million households, had at least one member age 12 or older who experienced one or more types of identity theft victimization.
- Among households in which at least one member experienced one or more types of identity theft, 64.1 percent experienced the misuse or attempted misuse of an existing credit card account in 2010.
- From 2005 to 2010, the percentage of all households with one or more type of identity theft that suffered no direct financial loss increased from 18.5 percent to 23.7 percent.[26]

A report from Javelin Strategy & Research covering identity fraud in 2012 included the following observation in the overview: "Identity fraud incidence increased in 2012 for the second consecutive year, affecting 5.26% of U.S. adults. This increase was driven by dramatic jumps in the two most severe fraud types, new account fraud (NAF) and account takeover fraud (ATF)."[27]

2.5 PHONE PHREAKING. Even in the earliest days of telephony, teenage boys played with the new technology to cause havoc. In the late 1870s, the new AT&T system in America had to stop using teenagers as switchboard operators:

> The boys were openly rude to customers. They talked back to subscribers, saucing off, uttering facetious remarks, and generally giving lip. The rascals took Saint Patrick's Day off without permission. And worst of all they played clever tricks with the switchboard plugs: disconnecting calls, crossing lines so that customers found themselves talking to strangers, and so forth.
>
> This combination of power, technical mastery, and effective anonymity seemed to act like catnip on teenage boys.[28]

2.5.1 2600 Hz. In the late 1950s, AT&T began switching its telephone networks to direct-dial long distance, using specific frequency tones to communicate among its switches. Around 1957, a blind seven-year-old child named Josef Engressia with perfect pitch and an emotional fixation on telephones learned to whistle the 2600-Hz pitch that interrupted long-distance telephone calls and allowed him to place a free long-distance call to anywhere in the world.[29] This emotionally disturbed person eventually renamed himself "Joybubbles" and is often described as the founder of phone phreaking—the manipulation of the phone system for unauthorized access to services.

John Draper was in the U.S. Air Force in 1964 when he began helping his colleagues place free phone calls. At the suggestion of Joybubbles, he used the whistles in Cap'n Crunch cereal boxes to generate the 2600-Hz tone and then, calling himself Captain Crunch, went on to create electronic tone synthesizers called *blue boxes*.[30] In the 1970s, Apple founders Steve Wozniak and Steve Jobs built blue boxes and, using the devices, perpetrated such pranks as calling the Vatican while pretending to be Henry Kissinger.[31]

A significant contributor to the growth of phreaking in the 1970s was the publication in 1971 of an article about phreaking in *Esquire Magazine,* which attracted the attention of many young technophiles.[32]

2.5.2 1982–1991: Kevin Poulsen. As the phone system shifted to greater reliance on computers, the border between phreaking and hacking began to blur. One of the important names from the 1980s period of fascination with everything phone-related was Kevin Poulsen.

Kevin Poulsen's autobiographical sketch is shown next.

> Kevin Poulsen first gained notoriety in 1982, when the Los Angeles County District Attorney's Office raided him for gaining unauthorized access to a dozen computers on the ARPANET, the forerunner of the modern Internet. Seventeen years old at the time, he was not charged, and went on to work as a programmer and computer security supervisor for SRI International in Menlo Park, California, then as a network administrator at Sun Microsystems.
>
> In 1987, Pacific Bell security agents discovered that Poulsen and his friends had been penetrating telephone company computers and buildings. After learning that Poulsen had also worked for a defense contractor where he'd held a SECRET level security clearance, the FBI began building an espionage case against the hacker.

Confronted with the prospect of being held without bail, Poulsen became a fugitive. While on the run, he obtained information on the FBI's electronic surveillance methods, and supported himself by hacking into Pacific Bell computers to cheat at radio-station phone-in contests, winning a vacation to Hawaii and a Porsche 944-S2 Cabriolet in the process.

After surviving two appearances on NBC's *Unsolved Mysteries*, Poulsen was finally captured on April 10th, 1991, in a Van Nuys grocery store, by a Pacific Bell security agent acting on an informant's tip. On December 4th, 1992, Poulsen became the first hacker to be indicted under U.S. espionage laws when the Justice Department charged him with stealing classified information. (18 U.S.C. 793).

Poulsen was held without bail while he vigorously fought the espionage charge. The charge was dismissed on March 18th, 1996.

Poulsen served five years, two months, on a 71-month sentence for the crimes he committed as a fugitive, and the phone hacking that began his case. He was freed June 4th, 1996, and began a three-year period of supervised release, barred from owning a computer for the first year, and banned from the Internet for the next year and a half.

Since his release, Poulsen has appeared on MSNBC, and on ABC's Nightline, and he was the subject of Jon Littman's flawed book, "The Watchman—the Twisted Life and Crimes of Serial Hacker Kevin Poulsen." His case has earned mention in several computer security and infowar tracts—most of which still report that he broke into military computers and stole classified documents.[33]

After his release from prison, Kevin Poulsen turned to journalism. He became an editor for *SecurityFocus* and then was hired as a senior editor at *Wired News.* He is a serious investigative reporter (e.g., he broke the story of sexual predators in MySpace)[34] and a frequent contributor to the "Threat Level" blog.[35]

2.6 DATA DIDDLING. One of the most common forms of computer crime since the start of electronic data processing is *data diddling*—illegal or unauthorized data alteration. These changes can occur before and during data input, or before output. Data-diddling cases have included bank records, payrolls, inventory data, credit records, school transcripts, telephone switch configurations, and virtually all other applications of data processing.

2.6.1 Equity Funding Fraud (1964–1973). One of the classic early data-diddling frauds was the Equity Funding case, which began with computer problems at the Equity Funding Corporation of America, a publicly traded and highly successful firm with a bright idea. The idea was that investors would buy insurance policies from the company and also invest in mutual funds at the same time, with profits to be redistributed to clients and to stockholders. Through the late 1960s, Equity's shares rose dizzyingly in price, and there were news magazine stories about this wunderkind of the Los Angeles business community.

The computer problems occurred just before the close of the financial year in 1964. An annual report was about to be printed, yet the final figures simply could not be extracted from the mainframe. In despair, the head of data processing told the president the bad news; the report would have to be delayed. Nonsense, said the president expansively (in the movie, anyway); simply make up the bottom line to show about $10 million in profits and calculate the other figures so it would come out that way. With trepidation, the DP chief obliged. He seemed to rationalize it with the thought that it was just a temporary expedient, and could be put to rights later in the real financial books.

The expected profit did not materialize, and some months later, it occurred to the executives at Equity that they could keep the stock price high by manufacturing false insurance policies that would make the company look good to investors. They therefore began inserting false information about nonexistent policyholders into the computerized records used to calculate the financial health of Equity.

In time, Equity's corporate staff got even greedier. Not content with jacking up the price of their stock, they decided to sell the policies to other insurance companies via the redistribution system known as reinsurance. Reinsurance companies pay money for policies they buy and spread the risk by selling parts of the liability to other insurance companies. At the end of the first year, the issuing insurance companies have to pay the reinsurers part of the premiums paid in by the policyholders. So in the first year, selling imaginary policies to the reinsurers brought in large amounts of real cash. However, when the premiums came due, the Equity crew "killed" imaginary policyholders with heart attacks, car accidents, and, in one memorable case, cancer of the uterus—in a male imaginary policyholder.

By late 1972, the head of DP calculated that by the end of the decade, at this rate, Equity Funding would have insured the entire population of the world. Its assets would surpass the gross national product of the planet. The president merely insisted that this showed how well the company was doing.

The scheme fell apart when an angry operator who had to work overtime told the authorities about shenanigans at Equity. Rumors spread throughout Wall Street and the insurance industry. Within days, the Securities and Exchange Commission had informed the California Insurance Department that they had received information about the ultimate form of data diddling: Tapes were being erased. The officers of the company were arrested, tried, and condemned to prison terms.[36]

2.6.2 1994: Vladimir Levin and the Citibank Heist. In February 1998, Vladimir Levin was sentenced to three years in prison by a court in New York City. Levin masterminded a major conspiracy in 1994 in which the gang illegally transferred $12 million in assets from Citibank to a number of international bank accounts. The crime was spotted after the first $400,000 was stolen in July 1994, and Citibank cooperated with the FBI and Interpol to track down the criminals. Levin was ordered to pay back $240,000, the amount he actually managed to withdraw before he was arrested.[37] The incident led to Citibank's hiring of Stephen R. Katz as the banking industry's first chief information security officer (CISO).

2.7 SALAMI FRAUD. In the salami technique, criminals steal money or resources a bit at a time. Two different etymologies are circulating about the origins of this term. One school of security specialists claim that it refers to slicing the data thin—like a salami. Others argue that it means building up a significant object or amount from tiny scraps—like a salami.

There were documented cases of salami frauds in the 1970s and 1980s, but one of the more striking incidents came to light in January 1993, when four executives of a Value Rent-a-Car franchise in Florida were charged with defrauding at least 47,000 customers using a salami technique. The federal grand jury in Fort Lauderdale claimed that the defendants modified a computer billing program to add five extra gallons to the actual gas tank capacity of their vehicles. From 1988 through 1991, every customer who returned a car without topping it off ended up paying inflated rates for an inflated

total of gasoline. The thefts ranged from $2 to $15 per customer—rather thick slices of salami but nonetheless difficult for the victims to detect.

Unfortunately, salami attacks are *designed* to be difficult to detect. The only hope is that random audits, especially of financial data, will pick up a pattern of discrepancies and lead to discovery. As any accountant will warn, even a tiny error must be tracked down, since it may indicate a much larger problem. For example, Cliff Stoll's famous adventures tracking down spies in the Internet began with an unexplained $0.75 discrepancy between two different resource accounting systems on UNIX computers at the Keck Observatory of the Lawrence Berkeley Laboratories. Stoll's determination to understand how the problem could have occurred revealed an unknown user; investigation led to the discovery that resource-accounting records were being modified to remove evidence of system use. The rest of the story is told in Clifford Stoll's book *The Cuckoo's Egg.*

2.8 LOGIC BOMBS. A logic bomb is a program that has deliberately been written or modified to produce results when certain conditions are met that are unexpected and unauthorized by legitimate users or owners of the software. Logic bombs may be within standalone programs, or they may be part of worms (programs that hide their existence and spread copies of themselves within a computer systems and through networks) or viruses (programs or code segments which hide within other programs and spread copies of themselves).

Time bombs are a subclass of logic bombs that "explode" at a certain time.

According to a National Security Council employee, the United States Government authorized insertion of a time bomb in software to control the Trans-Siberian natural gas pipeline that they knew would be stolen from U.S. sources by the Soviet government. "The result was the most monumental non-nuclear explosion and fire ever seen from space," said Thomas C. Reed.[38]

The infamous Jerusalem virus (also known as the Friday the 13th virus) of 1988 was a time bomb. It duplicated itself every Friday and on the thirteenth of the month, causing system slowdown; on every Friday the 13th after May 13, 1988, it also corrupted all available disks on the infected systems.

Other examples of notorious time bombs include:

- A common PC virus from the 1980s, *Cascade,* made all the characters fall to the last row of the display during the last three months of every year.
- The Michelangelo virus of 1992 was designed to damage hard disk directories on the sixth of March every year.
- In 1992, computer programmer Michael Lauffenburger was fined $5,000 for leaving a logic bomb at General Dynamics. His intention was to return after his program had erased critical data and be paid to fix the problem.[39]

The most famous time bomb of recent years was the Y2K (year 2000) problem. In brief, old programs used two-digit year codes that were based on the assumption that they applied to the twentieth century. As the twenty-first century approached, analysts warned of catastrophic consequences if the programs were not corrected to use four-digit years or otherwise adapt to the change of century.[40] In the event, the corrective measures worked and there were no disasters. Later analysis showed a positive correlation between investments in Y2K remediation and later profitability.[41]

2.9 EXTORTION. Computer data can be held for ransom. For example, according to Whiteside, in 1971, two reels of magnetic tape belonging to a branch of the Bank of America were stolen at Los Angeles International Airport. The thieves demanded money for their return. The owners ignored the threat of destruction because they had adequate backup copies.

Other early cases of extortion involving computers:

- In 1973, a West German computer operator stole 22 tapes and received $200,000 for their return. The victim did not have adequate backups.
- In 1977, a programmer in the Rotterdam offices of Imperial Chemical Industries, Ltd. (ICI) stole all his employer's tapes, including backups. Luckily, ICI informed Interpol of the extortion attempt. As a result of the company's forthrightness, the thief and an accomplice were arrested in London by officers from Scotland Yard.

In the 1990s, one of the most notorious cases of extortion was the 1999 theft of 300,000 records of customer credit cards from the CD Universe Web site by "Maxus," a 19-year-old Russian. He sent an extortion note that read: "Pay me $100,000 and I'll fix your bugs and forget about your shop forever … or I'll sell your cards [customer credit data] and tell about this incident in news." Refused by CD Universe owners, he promptly released 25,000 credit card numbers via a Web site that became so popular with criminals that Maxus had to limit access to one stolen number per visit.

2.10 TROJAN HORSES. Trojans are programs that pretend to be useful but that also contain harmful code or are just plain harmful.

2.10.1 1988 Flu-Shot Hoax. One of the nastiest tricks played on the shell-shocked world of early microcomputer users was the FLU-SHOT-4 incident of March 1988. With the publicity given to damage caused by destructive, self-replicating virus programs distributed through electronic bulletin board systems (BBSs), it seemed natural that public-spirited programmers would rise to the challenge and provide protective screening.

Flu-Shot-3 was a useful program for detecting viruses. Flu-Shot-4 appeared on BBSs and looked just like version 3; however, it actually destroyed critical areas of hard disks and any floppies present when the program was run. The instructions that caused the damage were not present in the program file until it was running; this self-modifying code technique makes it especially difficult to identify Trojans by simple inspection of the assembler-level code.

2.10.2 Scrambler, 12-Tricks, and PC Cyborg. Other early and notorious PC Trojans from the late 1980s that are still remembered in the industry included:

- The Scrambler (also known as the KEYBGR Trojan), which pretended to be a keyboard driver (KEYBGR.COM), but actually made a smiley face move randomly around the screen

- The 12-Tricks Trojan, which masqueraded as CORETEST.COM, a program for testing the speed of a hard disk, but actually caused 12 different kinds of damage (e.g., garbling printer output, slowing screen displays, and formatting the hard disk).
- The PC Cyborg Trojan (or "AIDS Trojan"), which claimed to be an AIDS information program but actually encrypted all directory entries, filled up the entire C disk, and simulated COMMAND.COM, but produced an error message in response to nearly all commands.

2.10.3 1994: Datacomp Hardware Trojan. On November 8, 1994, a correspondent reported to the *RISKS Forum Digest* that he had been victimized by a curious kind of Trojan:

> I recently purchased an Apple Macintosh computer at a "computer superstore," as separate components—the Apple CPU, and Apple monitor, and a third-party keyboard billed as coming from a company called Sicon.
>
> This past weekend, while trying to get some text-editing work done, I had to leave the computer alone for a while. Upon returning, I found to my horror that the text "welcome datacomp" had been inserted into the text I was editing. I was certain that I hadn't typed it, and my wife verified that she hadn't, either. A quick survey showed that the "clipboard" (the repository for information being manipulated via cut/paste operations) wasn't the source of the offending text.
>
> As usual, the initial reaction was to suspect a virus. Disinfectant, a leading anti-viral application for Macintoshes, gave the system a clean bill of health; furthermore, its descriptions of the known viruses (as of Disinfectant version 3.5, the latest release) did not mention any symptoms similar to my experiences.
>
> I restarted the system in a fully minimal configuration, launched an editor, and waited. Sure enough, after a (rather long) wait, the text "welcome datacomp" once again appeared, all at once, on its own.

Further investigation revealed that someone had put unauthorized code in the ROM chip used in several brands of keyboard. The only solution was to replace the keyboard. Readers will understand the possible consequences of a keyboard that inserts unauthorized text into, say, source code. Winn Schwartau, the renowned computer security expert, has coined the word "chipping" to refer to such unauthorized modification of firmware.

2.10.4 Keylogger Trojans. By the mid-2000s, software and hardware Trojans designed to capture logs of keystrokes and sometimes to transmit those logs via covert Internet connections had become a well-known tool of industrial espionage. The United States Department of Homeland Security issued a warning in December 2005 that included this overview:

> According to industry security experts, the biggest security vulnerability facing computer users and networks is email with concealed Trojan Horse software—destructive programs that masquerade as benign applications and embedded links to ostensibly innocent websites that download malicious code. While firewall architecture blocks direct attacks, email provides a vulnerable route into an organization's internal network through which attackers can destroy or steal information.
>
> Attackers try to circumvent technical blocks to the installation of malicious code by using social engineering—getting computer users to unwittingly take actions that allow the code to be installed and organization data to be compromised.

> The techniques attackers use to install Trojan Horse programs through email are widely available, and include forging sender identification, using deceptive subject lines, and embedding malicious code in email attachments.
>
> Developments in thumb-sized portable storage devices and the emergence of sophisticated keystroke logging software and devices make it easy for attackers to discover and steal massive amounts of information surreptitiously.[42]

2.10.5 Haephrati Trojan. A case that made the news in the mid-2000s began when Israeli author Amon Jackont was upset to find parts of the manuscript on which he was working posted on the Internet. Then someone tried to steal money from his bank account. Suspicion fell on his stepdaughter's ex-husband, Michael Haephrati. Police discovered a keystroke logger on Jackont's computer. It turned out that Haephrati had also sold spy software to clients; the Trojan was concealed in what appeared to be confidential email. Once installed on the victims' computers, the software sent surveillance data to a server in London, England.

Haephrati was detained by U.K. police and investigations began in Germany and Israel. Twelve people were detailed in Israel; eight others were under house arrest. Suspects included private investigators and top executives from industrial firms. Victims included Hewlett-Packard, Ace hardware stores, and a cable-communications company.

Michael and Ruth Haephrati were extradited from Britain for trial in Israel on January 31, 2006. They were accused of installing the Trojan horse program that activated a key logger with remote-reporting capabilities.[43]

In March 2006, the couple were indicted in Tel Aviv for corporate espionage.[44] They pleaded guilty to the charges[45] and were sentenced to four and two years of jail, respectively, as well as punished with fines.[46]

The story did not end there, however. Two years later, "Four members of the Israeli Modi'in Ezrahi private investigation firm were sentenced on Monday after they were found guilty of using Trojan malware to steal commercially sensitive information from their clients' competitors."[47] The report continues:

> Asaf Zlotovsky, a manager at the Modi'in Ezrahi detective firm, was jailed for 19 months. Two other employees, Haim Zissman and Ron Barhoum, were sent to prison for 18 and nine months respectively. The firm's former chief exec, Yitzhak Rett, the victim of an apparent accident when he fell down a stairwell during a break in police questioning back in 2005, escaped a jail sentence under a plea bargaining agreement. Rett was fined 250,000 Israeli Shekels (£36,500) and ordered to serve ten months' probation over his involvement in the scam.

However, an article in April 2008 reported that Michael Haephrati "claimed that there was no jail time, and that he was completely free. As a matter of fact he was going to continue to offer his Trojan Horse service but this time he would only work with 'law enforcement agencies.' "[48]

2.10.6 Hardware Trojans and Information Warfare. In the late 2000s, a flurry of news stories discussed the dangers of growing reliance on Chinese-manufactured computing components.

> U.S. Defense Department sources say privately that the level of Chinese cyberattacks obliges them to avoid Chinese-origin hardware and software in all classified systems and as many unclassified systems as fiscally possible. The high threat of Chinese cyberpenetrations into U.S. defense networks will be magnified as the Pentagon increasingly loses domestic sources of "trusted and classified" microchips.[49]

The discovery of counterfeit Cisco routers worsened concerns about the reliability of Chinese-manufactured network equipment.[50] The FBI, Immigration and Customs Enforcement (ICE), Customs and Border Protection (CBP), and the Royal Canadian Mounted Police (RCMP) worked together to track a massive pattern of counterfeit network hardware including Cisco routers; these investigations and seizures raised questions about the reliability and trustworthiness of such equipment, much of which was manufactured in the People's Republic of China. Although Cisco scientists examined some of the counterfeit equipment and found no back doors, concern was serious enough that government agencies created test chips to challenge quality assurance processes at military contractors:

> In April [2008], the Defense Advanced Research Projects Agency, part of the Defense Department, began distributing chips with hidden Trojan horse circuitry to military contractors participating in an agency program, Trusted Integrated Circuits. The goal is to test forensic techniques for finding hidden electronic trap doors, which can be maddeningly elusive. The agency is not yet ready to announce the results of the test, said Jan Walker, a spokeswoman for the agency.[51]

A 2011 report on hearings before the U.S. House Oversight and Government Reform Committee about the issue of Trojan backdoors in imported software and hardware included this assertion with references for each topic:

> … [E]mbedded malware lurking in consumer tech is not a new development. Since it's been happening for years and is hardly a national security secret, it's unclear why Schaffer hesitated so long before answering. There have been many incidents of malware-infected products being shipped to consumers, from hardware, to software, and even tainted peripheral devices. Malware has been sent pre-loaded in products like USBs, microchips, cameras, battery chargers, digital photo frames, webcams, printers, cell phones, motherboards or system boards, and hard drives.[52]

2.11 NOTORIOUS WORMS AND VIRUSES. The next sections briefly describe some of the outstanding incidents that are often mentioned in discussions of the history of malware.[53]

2.11.1 1970–1990: Early Malware Outbreaks. The ARPANET was the precursor of the Internet.[54] According to several reports:

> Sometime in the early 1970s, the Creeper virus was detected on ARPANET, a US military computer network which was the forerunner of the modern Internet. Written for the then-popular Tenex operating system, this program was able to gain access independently through a modem and copy itself to the remote system. Infected systems displayed the message, "I'M THE CREEPER: CATCH ME IF YOU CAN."
>
> Shortly thereafter, the Reaper program was anonymously created to delete Creeper. Reaper was a virus: it spread to networked machines and if it located a Creeper virus, Reaper would delete it. Even the participants are unable to say whether Reaper was a response to Creeper, or if it was created by the same person or persons who created Creeper in order to correct their mistake.[55]

By 1981, the Apple II computer was a popular system among hobbyists; the Elk Cloner virus spread via infected floppy disks and is regarded as "the first large-scale computer virus outbreak in history."[56]

In 1986, the Brain boot-sector virus was the first IBM-PCs malware to spread around the world. It was created by two brothers from Lahore, Pakistan, and included this text:

> Welcome to the Dungeon (c) 1986 Brain & Amjads (pvt) Ltd VIRUS_SHOE RECORD V9.0 Dedicated to the dynamic memories of millions of viruses who are no longer with us today - Thanks GOODNESS!! BEWARE OF THE er…VIRUS: this program is catching program follows after these messages….$#@%@!!

The Lehigh Virus appeared at Lehigh University in Pennsylvania in 1987 and damaged the files of several professors and students. This early program-infector targeted only *command.com* and was therefore extremely limited in its spread.

In 1988, the Jerusalem virus, a file infector that reproduced by inserting its code into EXE and COM files, caused a global PC epidemic.

Another noteworthy infection of 19988 came from the self-encrypting Cascade virus of 1988, which confused many naive users who interpreted the falling symbols on their screen as part of an unexpected screen saver. This virus was one of the earliest examples of the attempts to counter signature-based antivirus products.

2.11.2 December 1987: Christmas Tree Worm. In December 1987, users of IBM mainframe computers connected to the European Academic Research Network (EARN), BITNET, and the IBM company VNET were flooded with email bearing a character-based representation of a Christmas tree. A student at Technische Universität Clausthal[57] in Germany launched "a worm, written in an IBM-specific language called REXX."[58] The worm used the victim's list of correspondents to send copies of itself to everyone on the list.[59]

2.11.3 November 2, 1988: Morris Worm. On November 2, 1988, the Internet was rocked by the explosive appearance of unauthorized code on systems all over the world. At 17:00 EST on November 2, 1988, Robert T. Morris, a student at Cornell University in Ithaca, New York, released a worm into the Internet. By midnight, it had attacked VAX computers running 4 BSD UNIX and SUN Microsystems Sun 3 computers throughout the United States. One of the most interesting aspects of the worm's progress through the Internet was the almost complete independence of its path from normal geographical constraints. It sometimes leaped from coast to coast faster than it reached physically neighboring computer systems. The worm graphically demonstrated that cyberspace has its own geography.

The worm often superinfected its hosts, leading to slowdowns in overall processing speed. The first Internet warning ("We are under attack") was posted at 02:38 on November 3 to the TCP-IP list by a scientist at University of California at Berkeley. At 03:34, Andy Sudduth, a friend of Morris's at Harvard, posted a warning message ("There may be a virus loose on the internet") anonymously and included a few comments on how to stop the worm. Unfortunately, Spafford writes, the Internet was so severely impeded by the worm that this message was not widely distributed for over 24 hours.

By 6:00 on the morning of November 3, messages were creeping through the Internet with details of how the worm worked. The news spread via news groups such as the TCP-IP list, Usenix 4bsd-ucb-fixes, and the Usenet news.announce.important group. Spafford and his friends and colleagues on the Internet collaborated feverishly on providing patches against the worm.

Meanwhile, as word spread of the attack, some systems administrators began cutting their networks out of the Internet. The Defense Communications Agency isolated its Milnet and Arpanet networks from each other around 11:30 on November 3. At noon, machines in the science and technology center at the Stanford Research Institute were shut down.

By late on November 4, a comprehensive set of patches was posted on the Internet to defend systems against the worm. That evening, a *New York Times* reporter told Spafford that the author of the worm had been found.

By November 8, the Internet seemed to be back to normal. A group of concerned computer scientists met at the National Computer Security Center to study the incident and think about preventing recurrences of such attacks. Spafford put the incident into perspective with the comment that the affected systems were no more than 5 percent of the hosts on the Internet. It would be foolish to dismiss Morris's electronic vandalism as a prank or to claim that the worm alerted managers to weak security on their systems. Nonetheless, it is true that the incident contributed to the establishment of the Computer Emergency Response Team at the Software Engineering Institute of Carnegie-Mellon University. For these blessings, however, we owe no gratitude to Robert T. Morris.

In 1990, Morris was found guilty under the Computer Fraud and Abuse Act of 1986. The maximum penalties included five years in prison, a $250,000 fine, and restitution costs. Morris was ordered to perform 400 hours of community service, sentenced to three years probation, and required to pay $10,000 in fines. He was expelled from Cornell University.

His lawyers appealed the conviction to the Supreme Court of the United States. Their arguments included lack of evil intent (he did not mean to cause harm, honest—even though his worm took extraordinary precautions to conceal itself) and they deplored the scandalous behavior of Cornell University authorities, who had the temerity to search their own electronic mail message system to locate evidence that incriminated Morris. The lawyers also argued that sending a mail message might become a crime if Morris's conviction were upheld.

The Supreme Court upheld the decision by declining to hear the appeal.[60]

Robert T. Morris eventually became an associate professor in the Electrical Engineering and Computer Science Department of the Massachusetts Institute of Technology and a member of the Computer Science and Artificial Intelligence Laboratory.[61]

2.11.4 Malware in the 1990s. The most significant malware development of the 1990s was the release in July 1995 of the world's first widely distributed macro-language virus. The *macro.concept* virus made its appearance in MS-Word for Windows documents. It demonstrated how to use the macro programming language, common to many Microsoft products, to generate self-reproducing macros that spread from document to document. Within a few months, clearly destructive versions of this demonstration virus appeared.

Macro viruses were a dangerous new development. As explained in a recent history of viruses and antiviruses:

- Putting self-reproducing code in easily and frequently exchanged files, such as documents, greatly increased the infectiousness of the viruses.
- Virus writers shifted their attention to a much easier programming language than assembly.

- Email exchanges of infected documents were a far more effective mechanism for virus infection than exchanges of infected programs or disks.
- "[M]acro viruses were neither platform-specific, nor OS-specific. They were application-based."[62]

In the latter half of the 1990s, macro viruses replaced boot sector viruses and file infector viruses as a major type of malicious self-reproducing malware; during that period, additional types of script-based, network worms also increased.

Exhibit 2.1 shows the rise and fall of prevalence of macro viruses over the decade from discovery to extinction using data from the WildList archives. The WildList shows malware identified on user systems by at least two virus researchers.[63]

Roger Thompson summarizes the developments in malware in the 1990s in this way:

> By around 2000, macro viruses ceased to be a problem because the new version of MS-Office 2000 included features that blocked macro viruses. The next step in the evolution of malware was the mass mailers like the ILOVEYOU worm and then the network worms. These were easy to write and easy to obfuscate by varying the text contents, thus defeating signature scanners. These worms spread very quickly until the release of Windows XP Service Pack 2, which forced the Windows Firewall to be on by default. After that extinction-level event, criminals moved onward to creating mass mailers and bots which could spread malware and spam or cause distributed denial-of-service through communication via the trusted Web sites accessed through browsers that created a tunnel through the firewall.[64]

2.11.5 March 1999: Melissa. On Friday, March 26, 1999, the CERT/CC received initial reports of a fast-spreading new MS-Word macro virus. "Melissa" was written to infect such documents; once loaded, it uses the victim's MAPI-standard email address book to send copies of itself to the first 50 people on the list. The virus attaches an infected document to an email message with subject line "Subject: Important Message From <name>" where <name> is that of the inadvertent sender. The email message reads: "Here is that document you asked for … don't show anyone else;-)" and includes a MS-Word file as an infected attachment. The original infected document, "list.doc," was a compilation of URLs for pornographic Websites. However, as the virus spread, it was capable of sending any other infected document created by the victim.

Because of this high replication rate, the virus spread faster than any previous virus in history. On many corporate systems, the rapid rate of internal replication saturated email servers with outbound automated junk email. Initial estimates were in the range of 100,000 downed systems. Antivirus companies rallied immediately, and updates for all the standard products were available within hours of the first notices from CERT/CC.

The search for the originator of the Melissa email computer virus/worm began immediately after the outbreak. Initial findings traced the virus to Access Orlando, a Florida Internet Service Provider (ISP), whose servers were shut down by order of the FBI for forensic examination; the systems were then confiscated. That occurrence was then traced back to Source of Kaos, a free-speech Website where the virus may have lain dormant for months in a closed but not deleted virus-distributor's pages. Investigators discovered a serial number in the vector document, written with MS-Word; the undocumented serial number helped law enforcement when investigators circulated it on the Net to help track down the perpetrator.

The next steps turned to the value-added network AOL, where the virus was released to the public. The giant ISP's information helped to identify a possible suspect and

Exhibit 2.1 Rise and Fall in Macro Viruses in the WildList, 1996–2008

Year	Macro Viruses	Total Entries	Percentage Macro Virus
1996[a]	1	183	0.6%
1997[b]	27	239	11%
1998[c]	77	258	30%
1999[d]	46	129	36%
2000[e]	108	175	62%
2001[f]	145	228	64%
2002[g]	103	198	52%
2003[h]	68	205	33%
2004[i]	51	261	20%
2005[j]	22	399	6%
2006[k]	19	804	2%
2007[l]	5	797	0.6%
2008[m]	0	590	0.0%

[a]WildList Organization International, "PC Viruses in the Wild—January 10, 1996," www.wildlist.org/WildList199601.htm.
[b]WildList Organization International, "PC Viruses in the Wild—February, 1997," www.wildlist.org/WildList199702.htm.
[c]WildList Organization International, "PC Viruses in the Wild—January, 1998," www.wildlist.org/WildList199801.htm.
[d]WildList Organization International, "PC Viruses in the Wild—January 1999," www.wildlist.org/WildList199001.htm.
[e]WildList Organization International, "PC Viruses in the Wild—January, 2000," www.wildlist.org/WildList200001.htm.
[f]WildList Organization International, "PC Viruses in the Wild—January, 2001," www.wildlist.org/WildList200101.htm.
[g]WildList Organization International, "PC Viruses in the Wild—January, 2002," www.wildlist.org/WildList200201.htm.
[h]WildList Organization International, "PC Viruses in the Wild—January, 2003," www.wildlist.org/WildList200301.htm.
[i]WildList Organization International, "PC Viruses in the Wild—January, 2004," www.wildlist.org/WildList200401.htm.
[j]WildList Organization International, "PC Viruses in the Wild—January, 2005," www.wildlist.org/WildList200501.htm.
[k]WildList Organization International, "PC Viruses in the Wild—January, 2006," www.wildlist.org/WildList200601.htm.
[l]WildList Organization International, "PC Viruses in the Wild—January, 2007," www.wildlist.org/WildList200701.htm.
[m]WildList Organization International, "PC Viruses in the Wild—January, 2008," www.wildlist.org/WildList200801.htm.

by April 2, the FBI arrested David L. Smith (age 30) of Aberdeen, New Jersey. Smith apparently panicked when he heard the FBI was on the trail of the Melissa spawner and he threw away his computer—stupidly, into the trash at his own apartment building.

Smith was charged with second-degree offenses of interruption of public communication, conspiracy to commit the offense and attempt to commit the offense, third-degree theft of computer service, and third-degree damage or wrongful access to computer systems. If convicted, Smith faced a maximum penalty of $480,000 in fines and 40 years in prison. On December 10, 1999, Smith pleaded guilty to all federal charges and agreed to every particular of the indictment, including the estimates by the International Computer Security Association of at least $80 million of consequential damages due to the Melissa infections.[65]

2.11.6 May 2000: I Love You. Starting around May 4, 2000, email users opened messages from familiar correspondents with the subject line "I love you"; many then opened the attachment, LOVE-LETTER-FOR-YOU.txt.vbs, which infected the user's email address book and initiated mass mailing of itself to all the contacts. The "Love Bug" was the fastest-spreading worm to that time, infecting computers all over the world, starting in Asia, then Europe.[66]

On May 11, Filipino computer science student Onel de Guzman of AMA Computer College in Manila admitted to authorities that he may "accidentally have launched the destructive Love Bug virus out of youthful exuberance." He did not admit that he had created the malware himself; however, the name GRAMMERSoft appeared in the computer code of the virus, and that was the name of a computer group to which the 23-year-old de Guzman belonged.[67]

In September 2000, de Guzman participated in a live chat hosted by CNN.com; he vigorously defended virus-writing and blamed the creators of vulnerable systems for releasing poorly designed software. He refused to take responsibility for writing the worm.[68]

Philippine authorities tried to prosecute de Guzman but had to drop their attempts in August 2000 for lack of sufficient evidence. Due to the lack of computer crime laws at the time, it was impossible for other countries such as the United States to extradite the suspect: International principles of dual criminality require equivalent laws in both jurisdictions before extradition can proceed.

By October 2000, de Guzman had refused to take responsibility for writing the worm and publicly stated, "'I admit I create viruses, but I don't know if it's one of mine.... If the source code was given to me, I could look at it and see. Maybe it is somebody else's, or maybe it was stolen from me."[69]

The "I Love You" case was a wake-up call for the international community to think about standardizing computer crime laws around the globe.[70]

2.11.7 July 2010 Stuxnet. In July 2010, reports surfaced of a zero-day threat to SCADA systems using Siemens AG's Simatic WinCC and PCS 7 software. Analysts found that the Stuxnet worm was designed for industrial espionage; however, the same techniques could have been used for sabotage. Experts expressed concern that the worm was signed using valid digital certificates from Taiwanese companies and that the complex code implied considerable knowledge of the SCADA software.[71] Further analysis of the malware code suggested that the software was developed by the United States and Israel and used at least as early as November 2007.[72]

2.12 SPAM. Chapter 20 in this *Handbook* includes a detailed history of unsolicited commercial email and the reason it is called *spam*. This section looks solely at a seminal abuse of the USENET in 1994 and trends in spam over the next decade.

2.12.1 1994: Green Card Lottery Spam. On April 2, 1994, Laurence A. Canter and Marthas S. Siegel posted an advertisement for legal services connected to the U.S. Government's Green Card Lottery to over 6,000 USENET groups. Instead of cross-posting their commercial message, they used a script to post a copy of the message separately to every group. The former method would have shown the message to USENET users once; Canter and Siegel's abuse of the USENET made their ad show up in every affected group to which users subscribed.[73]

Reaction worldwide was massive. Automated cancelbots trolled the USENET deleting the unwanted messages; the attorneys' ISP was so overloaded with email complaints that its servers crashed. Canter and Siegel were reviled in postings and newspaper articles.[74] Their unsavory backgrounds were posted in discussion groups, including details of disciplinary hearings before the Florida Bar and accusations of dishonesty and unprofessional behavior.[75]

Unfazed, the couple published a book about how to abuse the Internet using spam and defended their actions in interviews as an expression of freedom of speech; they dismissed critics as "wild-eyed zealots" or as commercial interests intent on controlling the Internet for their own gain.[76]

Canter was eventually disbarred in Tennessee, in part for his spamming.[77] He remained unrepentant; in 2002, he spammed 50,000 K–12 teachers with an advertisement for a book whose title he liked so he could harvest payments for referrals from Amazon.[78]

2.12.2 Spam Goes Global. Over the next decade, the incidence of spam grew explosively. By 2007, spam watchers and anti-spam companies reported that around 88 percent of all email traffic on the Internet was spam. Spammers caused so much irritation that companies developed software and hardware solutions for filtering email by content. Spammers responded by increasing the number of images in their spam, making content filtering more difficult. At one point, the amount of spam grew 17 percent between one day and the next as spammers began pumping PDF files into spam pipelines.[79]

Botnets spawned through infected zombie machines established rogue SMTP nodes using innocent (and ignorant) PC users' computers and persistent high-speed Internet connections.[80] Spam currently provides a major vector for fraud by deceit, including in particular 4-1-9 advance fee fraud and phishing attacks.[81] Advance-fee fraud usually consists of enticements to participate in the theft of ill-gotten gains such as bank deposits belonging to dead people or stolen from poor countries; the dupes who agree to participate in such illegality are promised millions of dollars—only to be told that they suddenly have to send cash for unexpected bribes or fees. If they do so, they are asked for more … and more … and more. Phishing involves sending email messages that are supposed to look like official, usually alarming, warnings from banks and other institutions; victims click on links that look like one thing but actually go to the criminals' Websites. There the victims cheerfully type in their user identification, passwords, bank account numbers, and all manner of other confidential information useful for identity theft.[82] Advance-fee fraud and phishing are discussed in Chapter 20 in this *Handbook*.

2.13 DENIAL OF SERVICE. Denial of service results from exhaustion or destruction of necessary resources and is thoroughly discussed in Chapter 18. However, a couple of denial-of-service attackers stand out among all the others in the last two decades: the Unamailer and Mafiaboy.

2.13.1 1996: Unamailer. In August 1996, someone using the pseudonym "johnny [x]chaotic" claimed the blame for a massive mail-bombing run based on fraudulently subscribing dozens of victims to hundreds of mailing lists. The denial of service was the result in part of the naïveté of list managers who accepted subscriptions for any email address from any other email address. In a rambling and incoherent letter posted on the Net, (s)he made rude remarks about famous and not-so-famous

people, whose capacity to receive meaningful email was then obliterated by up to thousands of unwanted messages a day.[83] "The first attack, in August, targeted more than 40 individuals, including Bill Clinton and Newt Gingrich and brought a torrent of complaints from the people who found their names sent as subscribers to some 3,000 E-mail lists."[84]

Someone claiming to be the same "Unamailer" (as the news media labeled him or her in reference to the Unabomber) launched a similar mass-subscription mail-bombing run in late December.

> This attack is estimated to involve 10,139 listservs groups, 3 times greater than the one that took place in the summer, also at xchaotic's instigation. If each mailing list in this attack sent the targeted individuals just a modest 10 letters to the subscribers' computers those individuals would receive more than 100,000 messages. If each listing system sent 100 messages—and many do—then the total messages could tally 1,000,000.[85]

In December, the attacker(s) sneered at list administrators for failing to use authentication before allowing subscriptions and wrote that they would continue their attacks until practices changed.[86]

Partly as a result of the Unamailer's depredations, list administrators did in fact change their practices—not that anyone thanked Johnny [x]chaotic for his method of persuasion.

2.13.2 2000: MafiaBoy. On February 8, 2000, Yahoo.com suffered a three-hour flood from a distributed denial-of-service (DDoS) attack and lost its capacity to serve Web pages to visitors. The next day, the same technique was extended to Amazon.com, eBay.com, Buy.com, and CNN.com.[87] Later information also showed that Charles Schwab, the online stock brokerage, had been seriously impeded in serving its customers because of the DDoS. Buy.com managers were particularly disturbed because the attack occurred on the day of their initial public offering. As a result of the attacks, a number of firms formed a consortium to fight DDoS attacks.[88]

Investigation by the RCMP and the FBI located a 15-year-old child in west-end Montreal who used a modem to control zombies in his DDoS escapade:

> On April 15, 2000, the RCMP arrested a Canadian juvenile known as Mafiaboy for the February 8th DDoS attack on CNN in Atlanta, Georgia. On August 3, 2000, Mafiaboy was charged with 64 additional counts. On January 18, 2001, Mafiaboy appeared before the Montreal Youth Court in Canada and pleaded guilty to 56 counts. These counts included mischief to property in excess of $5,000 against Internet sites, including CNN.com, in relation to the February 2000 attacks. The other counts related to unauthorized access to several other Internet sites, including those of several US universities. On September 12, 2001, Mafiaboy appeared before the Montreal Youth Court in Canada and was sentenced to eight months "open custody," one year probation, and restricted use of the Internet.[89]

MafiaBoy's name was not released by Canadian authorities because of Canadian laws protecting juveniles, although several U.S. reporters distributed his identity in their publications. His chief contribution to the history of computer crime was to demonstrate asymmetric warfare in cyberspace.[90] His actions showed that even an ignorant child with little knowledge of computing could use low-tech hardware and tools available to anyone on the Internet to cripple major organizations.

2.14 HACKER UNDERGROUND. Newcomers to the field of information assurance will encounter references to the computer underground in texts, articles, and

discussions. The sections that follow provide thumbnail sketches of some of the key groups and events in the shadowy world of criminal hacking (known as *black hats,* in contrast to *white hats,* who are law enforcement and establishment security experts), and the intermediate range of well-intentioned rebels who use unorthodox means to challenge corporations and governments over what they see as security failings (these people are often called *gray hats*).

2.14.1 1981: Chaos Computer Club. On September 12, 1981, a group of German computer enthusiasts with a strong radical political orientation formed the Chaos Computer Club (CCC) in Hamburg.[91] One of their first achievements was to demonstrate a serious problem in the Bundespost's (German post office) new Bilschirmtext (BTX) interactive videotext service in 1984, not long after the service was announced.[92] The CCC used security flaws in BTX to transfer a sizable amount of money into their own bank account through a script that ran overnight as a demonstration to the press (returning the money publicly).

After the Legion of Underground (LoU) announced on January 1, 1999, that they would attack and disable the computer systems of the People's Republic of China and of Iraq, a coalition of hacker organizations including the CCC announced opposition to the move. "We strongly oppose any attempt to use the power of hacking to threaten or destroy the information infrastructure of a country, for any reason," the coalition said. "Declaring war against a country is the most irresponsible thing a hacker group could do. This has nothing to do with hacktivism or hacker ethics and is nothing a hacker could be proud of," the coalition said in the statement.

The CCC has, in general, challenged the general view that "hacker" necessarily means "criminal hacker."[93] Their annual Chaos Communications Conferences have proven to be a site of technology exchange and serious discussion of information security issues. Their continued commitment to the rule of law (except where their own activities are concerned), and their willingness to engage authorities in the courts when necessary has gained them an unusual degree of credibility and acceptance in the information security community as relatively pale-gray hats.[94]

2.14.2 1982: The 414s. One morning in June 1982, a system administrator for a DEC VAX 11/780 minicomputer at the Memorial Sloan-Kettering Cancer Center in Manhattan found his system down. Investigation led to the discovery that his and dozens of other systems around the country were being hacked by Milwaukee-area teenagers and others aged 15 to 22. The youths called themselves the 414s after the Milwaukee area code.

> Using home computers connected to ordinary telephone lines, they had been breaking into computers across the U.S. and Canada, including one at a bank in Los Angeles, another at a cement company in Montreal and, ominously, an unclassified computer at a nuclear weapons laboratory in Los Alamos, [New Mexico].[95]

In March 1984, "two members of Milwaukee's 414 Gang … pleaded guilty to misdemeanor charges of making obscene or harassing phone calls. Maximum sentence for each charge: six months in jail and a $500 fine."[96]

2.14.3 1984: Cult of the Dead Cow. Another influential criminal-hacker group is the Cult of the Dead Cow (cDc), which used to sport amusing (although intentionally offensive to some) cartoons such as that of a crucified cow.[97] The cDc

was noted for its consistent use of humor and parody; for example, "Swamp Rat's" 1985 article on building "The infamous … GERBIL FEED BOMB" included instructions such as "Light the fuse if you put one in. If you dropped a match into it, then go to the nearest phone, dial '911' and tell the nice people that you have a large number of glass shards embedded in your lower body. An ambulance should be there soon."[98]

The cDc became important proponents of hactivism in the 1990s—the use of criminal hacking techniques for political purposes. They also released a number of hacking tools, of which Back Orifice (BO) and especially Back Orifice 2000 (BO2K) were notorious examples. BO2K was ostensibly a remote administration tool but was in fact a Trojan that ran in stealth mode and allowed remote control of infected machines.[99] Some observers felt that presenting BO2K as a legitimate tool was another instance of cDc's satirical bent: The idea that anyone would consider software written by criminal hackers as a trustworthy administration tool struck them as ludicrous.

2.14.4 1984: *2600: The Hacker Quarterly*. Eric Corley founded *2600: The Hacker Quarterly* in 1984. This publication has become a standard-bearer for proponents of criminal hacking. The magazine has published a steady stream of explanations of how to exploit specific vulnerabilities in a wide range of operating systems and application environments. In addition, the editor's political philosophy has influenced more than one generation of black-hat and gray-hat hackers:

> In the worldview of *2600*, the tiny band of technocrat brothers (rarely, sisters) are a besieged vanguard of the truly free and honest. The rest of the world is a maelstrom of corporate crime and high-level governmental corruption, occasionally tempered with well-meaning ignorance. To read a few issues in a row is to enter a nightmare akin to Solzhenitsyn's, somewhat tempered by the fact that *2600* is often extremely funny.[100]

2.14.5 1984: Legion of Doom. The DC Comics empire created an animated cartoon series called *Super Friends* that appeared in 1973; it starred various DC Comics heroes, such as Superman, Aquaman, Wonder Woman, and Batman.[101] In a follow-up series called *Challenge of the Super Friends* that ran from 1978 through 1979, the archenemies of these heroes were a group known as the *Legion of Doom,* which included Lex Luthor, archenemy of Superman.[102] A group of phone phreakers who later turned to criminal hacking called themselves the Legion of Doom (LOD); their founder called himself "Lex Luthor." Another major member was Loyd Blankenship ("The Mentor").

Bruce Sterling describes the LOD as an influential hacker underground group of the 1980s and one of the earliest to capitalize on regular publication of their findings of vulnerabilities and exploits in the phone system and then in computer networks:

> LOD members seemed to have an instinctive understanding that the way to real power in the underground lay through covert publicity. LOD were flagrant. Not only was it one of the earliest groups, but the members took pains to widely distribute their illicit knowledge. Some LOD members, like "The Mentor," were close to evangelical about it. *Legion of Doom Technical Journal* began to show up on boards throughout the underground.
>
> *LOD Technical Journal* was named in cruel parody of the ancient and honored *AT&T Technical Journal*. The material in these two publications was quite similar—much of it, adopted from public journals and discussions in the telco community. And yet, the predatory attitude of LOD made even its most innocuous data seem deeply sinister; an outrage; a clear and present danger.[103]

In the later 1980s, the LOD actually helped law enforcement on occasion by restraining malicious hackers.

One of the best-known members was Chris Goggans, whose handle was "Erik Bloodaxe"; he was also an editor of *Phrack* and later became part of the Masters of Deception (MOD), which was involved in a conflict with LOD in 1990 and 1991 known in hacker circles as "The Great Hacker War."[104]

Another well-known hacker who started in LOD and moved to MOD was Mark Abene ("Phiber Optik"), who was eventually imprisoned for a year after pleading guilty in federal court to conspiracy and unauthorized access to federal-interest computers (a violation of 18 USC 1030(a), the Computer Fraud and Abuse Act of 1986).[105] Abene's punishment was the subject of much protest in the hacker community and elsewhere.[106]

2.14.6 1985: *Phrack*. *Phrack* began publishing in November 1985. With a new issue every month or two at first, the electronic magazine continued uninterrupted distribution of technical information and rants. The uncensored commentary provided a fascinating glimpse of some of the personalities and worldviews of its contributors and editors, including Taran King and Craig Neidorf (later to become famous as "Knight Lightning" and for his involvement in an abortive prosecution involving BellSouth documents). For example, *Phrack* published what became known as the "Hacker Manifesto"—held up by criminal hackers as a light unto the nations ("Written almost 15 years ago by The Mentor, this should be taped up next to everyone's monitor to remind them who we are, this rang true with Hackers, but it now rings truth to the internet generation."[107]), but viewed with skepticism by security professionals. It read in part:

> This is our world now … the world of the electron and the switch, the beauty of the baud. We make use of a service already existing without paying for what could be dirt-cheap if it wasn't run by profiteering gluttons, and you call us criminals. We explore … and you call us criminals. We seek after knowledge … and you call us criminals. We exist without skin color, without nationality, without religious bias … and you call us criminals. You build atomic bombs, you wage wars, you murder, cheat, and lie to us and try to make us believe it's for our own good, yet we're the criminals.
>
> Yes, I am a criminal. My crime is that of curiosity. My crime is that of judging people by what they say and think, not what they look like. My crime is that of outsmarting you, something that you will never forgive me for.
>
> I am a hacker, and this is my manifesto. You may stop this individual, but you can't stop us all … after all, we're all alike.[108]

In the 1990s, publication frequency faltered, falling to once every three to six months until the editors announced the final issue, #63, for August 2005. However, publication resumed under new editorial leadership in May 2007 with issue 64; given that issue 65 did not come out until April 2008, the magazine's heyday is presumably past.

2.14.7 1989: Masters of Deception. The Masters of Deception (MOD) were a New York hacker group active from about 1989 through 1992.[109] Among the most notorious criminal hackers in the group was "Phiber Optik" (Mark Abene, born in 1972), who was unusually visible in the media:

> Phiber Optik in particular was to seize the day in 1990. A devotee of the 2600 circle and stalwart of the New York hackers' group "Masters of Deception," Phiber Optik was a splendid exemplar of the computer intruder as committed dissident. The eighteen-year-old Optik, a high-school dropout and part-time computer repairman, was young, smart, and ruthlessly obsessive, a sharp-dressing, sharp-talking digital dude who was utterly and airily contemptuous of anyone's rules

but his own. By late 1991, Phiber Optik had appeared in *Harper's*, *Esquire*, the *New York Times*, in countless public debates and conventions, even on a television show hosted by Geraldo Rivera.[110]

2.14.8 1990: Operation Sundevil. After two years of investigation, on May 7, 8, and 9, 1990, 150 FBI agents, aided by state and local authorities, raided presumed criminal-hacker organizations allegedly involved in credit-card abuse and theft of telephone services. They seized 42 computers and 23,000 disks from locations in 14 cities. Targets were principally sites running discussion boards, some of which were classified as "hacker boards." However, two years after the raid, there were only three indictments (resulting in three guilty pleas). Evidence began to accumulate that much of the evidence seized in the raids was useless.[111] Bruce Sterling spent a year and a half researching the operation and concluded that it was largely a propaganda effort:

> … An unprecedented action of great ambition and size, Sundevil's motives can only be described as political. It was a public-relations effort, meant to pass certain messages, meant to make certain situations clear: both in the mind of the general public, and in the minds of various constituencies of the electronic community.
>
> First—and this motivation was vital—a "message" would be sent from law enforcement to the digital underground. This very message was recited in so many words by Garry M. Jenkins, the Assistant Director of the US Secret Service, at the Sundevil press conference in Phoenix on May 9, 1990, immediately after the raids. In brief, hackers were mistaken in their foolish belief that they could hide behind the "relative anonymity of their computer terminals." On the contrary, they should fully understand that state and federal cops were actively patrolling the beat in cyberspace—that they were on the watch everywhere, even in those sleazy and secretive dens of cybernetic vice, the underground boards.[112]

2.14.9 1990: Steve Jackson Games. Two months before the Operation Sundevil raids, but (contrary to popular conflation of the two) in a completely separate operation, a role-playing game company called Steve Jackson Games in Austin, Texas, was raided on March 1, 1990. The Secret Service seized computers and disks at the company's offices and also at the home of one of their employees, Loyd Blankenship—"The Mentor," formerly of the LOD. Blankenship was writing a role-playing game called GURPS Cyberpunk, which the agents interpreted as "a handbook for computer crime." Some of the equipment seized in the raid was returned four weeks later; most but not all was returned four months later. The company nearly went bankrupt as a result of the sequestration of critical resources.[113]

Outrage in the computing community spread beyond the underground. Mitch Kapor, John Barlow, and John Gilmore founded the Electronic Frontier Foundation in part because of their outrage over the treatment of Steve Jackson Games:

> … We got the attorneys involved, and then we asked them to look into what was going on with a variety of government investigations and prosecutions. We identified a couple of particular legal situations, like Craig Neidorf in Chicago and Steve Jackson Games, where there seemed to us to have been a substantial overstepping of bounds by the government and an infringement on rights of free speech and freedom of the press. We were in the process of deciding how to intervene when we also realized very clearly that we didn't want to be a legal defense fund as that was too narrow. What was really needed was to somehow improve the discourse about how technology is going to be used by society; we need to do things in the area of public education and policy development.[114]

Steve Jackson Games sued the Secret Service for damages and were awarded $50,000 in damages and more than $25,000 in attorney's fees.[115] The case had a lasting effect

on how law enforcement officials carried investigations of computer crimes and seizure of electronic evidence.

2.14.10 1992: L0pht Heavy Industries. In 1992, a group of computer enthusiasts arranged to store their spare equipment in some rented space in Boston. They collaborated on analysis of vulnerabilities, especially Microsoft product vulnerabilities, and gained a reputation for contributing serious research to the field and for appearing at security conferences. Their "L0phtCrack" program was adopted by many system administrators for testing password files to locate easy-to-guess passwords; members even testified before a Senate Subcommittee on Government Cybersecurity in 1998 (saying they could take down the Internet in half an hour).[116] Famous handles from the group included "Brian Oblivion," "Kingpin," "Mudge," "Space Rogue," "Stefan von Neumann," "Tan," and "Weld Pond."[117]

The group caused ripples in both the underground and aboveground security communities when their company, L0pht Heavy Industries, was purchased by security services firm @stake, Inc. in 2000. @stake was eventually bought by Symantec in 1994.[118]

2.14.11 2004: Shadowcrew. Stealing physical credit cards and creating fake ones are part of the criminal technique called "carding." One of the significant successful investigations and prosecutions of an international credit card fraud ring of the 2000 decade began with the U.S. Secret Service's *Operation Firewall* in late 2004. The investigators discovered a network of more than 4,000 members communicating through the Internet and conspiring to use phishing, spamming, forged identity documents (e.g., fake driver's licenses), creation of fake plastic credit cards, resale of gift cards bought with fake credit cards, fencing of stolen goods via eBay, and interstate or international funds transfers using electronic money such as E-Gold and Web Money.

In October 2004, the Department of Justice indicted 19 of the leaders of Shadowcrew.[119] By November 2005, 12 of these people had already pleaded guilty to charges of conspiracy and trafficking in stolen credit card numbers with losses of more than $4 million.[120]

In February 2006, Shadowcrew leader Kenneth J. Flury, 41, of Cleveland Ohio, was sentenced to 32 months in prison with three years of supervised release and $300,000 in restitution to Citibank.[121] In June 2006, cofounder Andrew Mantovani, 24, of Scottsdale, Arizona, was fined $5,000 and also received 32 months of prison with three years of supervised release. Five other indicted Shadowcrew criminals were sentenced with him. By that time, a total of 18 of 28 indicted suspects had already pleaded guilty.[122]

2.14.12 Late 2000s: Russian Business Network (RBN). The Russian Business Network (RBN) may have originated as a legitimate Web hosting company in 2006:

> According to internet security company Verisign, which in June published an extensive investigation into the Russian outfit (tinyurl.com/ywvgpg), RBN was registered as an internet site in 2006.
>
> Initially, much of its activity was legitimate. But apparently the founders soon discovered that it was more profitable to host illegitimate activities and started hiring its services to criminals. Verisign says simply that it is now "entirely illegal." Since then its activities have been monitored by a number of organisations, including the London-based anti-spam group Spamhaus. "RBN is among the world's worst spammer, child-pornography, malware, phishing

and cybercrime hosting networks," says a spokesman. "It provides 'bulletproof' hosting, but is probably involved in the crime too."[123]

A researcher for the Internet Storm Center, "David Bizeul, spent the past three months researching the Russian Business Network (RBN). The RBN is a virtual safe house for Russian criminals responsible for malicious code attacks, phishing attacks, child pornography, and other illicit operations …" Bizeul's study is a 70-page report with extensive documentation about the criminal activities of the RBN.[124] The group has supported malware diffusion, spam, phishing, denial of service, distribution of cyberattack tools, pornography, and child pornography.

A 2011 report by David Goldman included the following useful insights:

"It's not like the Mafia, it is a Mafia running these operations," said Karim Hijazi, CEO of botnet monitoring company Unveillance. "The Russian Mafia are the most prolific cybercriminals in the world."

Organized cybercrime is a truly international affair, but the most advanced attacks tend to stem from Russia. The Russian mob is incredibly talented for a reason: After the Iron Curtain lifted in the 1990s, a number of ex-KGB cyberspies realized they could use their expert skills and training to make money off of the hacked information they had previously been retrieving for government espionage purposes.

Former spies grouped together to form the Russian Business Network, a criminal enterprise that is capable of some truly scary attacks. It's just one of many organized cybercriminal organizations, but it's one of the oldest and the largest.

"The Russians have everyone nailed cold in terms of technical ability," said Greg Hoglund, CEO of cybersecurity company HBGary. "The Russian crime guys have a ridiculous toolkit. They're targeting end users in many cases, so they have to be sophisticated."[125]

2.14.13 Anonymous. In 2003, political activists with a penchant for computer skills formed a loose association calling itself *Anonymous* for collaboration in a range of cyberattacks on targets its members disliked. The philosophy of the group explicitly rejects any centralized controls; anyone can claim to be a member of Anonymous.

In 2008, self-identified members of the movement labeling their efforts *Chanology*[126] attacked the Church of Scientology (readers interested in following the reference provided in the end note should be aware that the site is loaded with pornographic advertisements for pornography sites). Members also harassed organizations attempting to strengthen intellectual property laws and enforcement or antipiracy restrictions. Other targets of the nonorganization include the Epilepsy Foundation, hip-hop Websites, Sarah Palin's political campaign, the government of Iran, the government of Australia, and the Tea Party chapter in Oregon.

One of the most publicized campaigns was in support of Julian Assange, leader of the WikiLeaks Foundation, whose group made public more than a million documents classified by the United States and other governments as restricted or secret and revealing embarrassing details of several wars and internal communications among diplomats.

In January 2013, members announced that they would release large amounts of U.S. Government–restricted information. They let the world know about their plans by posting their messages on a hacked U.S. Government Website.[127]

2.14.14 2013: Unlimited Operations. In May 2013, eight criminal hackers, New York City area members of a much larger worldwide ring of cybercriminals calling themselves Unlimited Operations, were charged with theft of more than $45 million from automated teller machines (ATMs) around the planet. The gang "used

sophisticated intrusion techniques to hack into the systems of global financial institutions, steal prepaid debit card data, and eliminate withdrawal limits. The stolen card data was then disseminated worldwide and used in making fraudulent ATM withdrawals on a massive scale across the globe …"

In the first phase, the criminals broke into National Bank of Ras Al-Khaimah PSC (RAKBANK) in the United Arab Emirates. Using these compromised data, the criminal network completed more than 4,500 ATM transactions in 20 hours and stole more than $5 million.

The second phase began "… on the afternoon of February 19 and lasted into the early morning of February 20, 2013. This operation again breached the network of a credit card processor that serviced MasterCard prepaid debit cards, this time issued by the Bank of Muscat, located in Oman." Total losses from 36,000 transactions in 24 countries netted $40 million in cash from ATMs.[128]

2.15 INDUSTRIAL ESPIONAGE. Why spend money developing competitive products when you can steal the work once it's ready to apply? Many firms in countries with little or no rule of law have taken advantage of poor security, outsourcing, and liberal immigration policies to steal intellectual property and compete at a discount with the originators of the ideas.

- In 2001, Junsheng Wang of Bell Imaging Technologies pled guilty to violation of 18 USC 132(a)(2) by stealing trade secrets fom Acuson Corporation. The Counterintelligence News and Developments (CIND) report noted, "In pleading guilty, Wang admitted that prior to August 24, 2000, that he took without authorization and copied for Bell Imaging a document providing the architecture for the Sequoia ultrasound machine that contained the trade secrets of Acuson Corporation. According to Wang's plea agreement, he had been able to obtain access to the Acuson trade secret materials because his wife was employed as an engineer at that company and because she had brought that document into their home. After he had copied the document, he took it with him on business trips to the People's Republic of China, turning it over to Bell Imaging during 2000."[129]
- In May 2001, Federal authorities arrested two Lucent scientists and a third man described as their business partner on May 4, charging them with stealing source code for software associated with Lucent's PathStar Access Server and sharing it with Datang Telecom Technology Co., a Beijing firm majority-owned by the Chinese government. The software is considered a "crown jewel" of the company. Chinese nationals Hai Lin and Kai Xu were regarded as "distinguished members" of Lucent's staff up until their arrests. The motivation for the theft, according to court documents, was to build a networking powerhouse akin to the "Cisco of China." The men faced charges of conspiracy to commit wire fraud, punishable by a maximum five years in prison and a $250,000 fine.[130] In April 2002, the two were also charged with stealing secrets from four companies in addition to Lucent: Telenetworks, NetPlane Systems, Hughes Software Systems, and Ziatech. An additional Chinese national, Yong-Qing Cheng was also charged. They developed a joint venture with the Datang Telecom Technology Company of Beijing to sell a clone of Lucent's Path Star data and voice transmission system to Internet providers in China.[131]
- In September 2002, the 3DGeo company in Mountain View, CA accused Shan Yanming, an employee of the China National Petroleum Corporation on loan to

the company, of industrial espionage for trying to steal the software designed for using seismic data to map oil deposits. He was caught trying to download corporate data to his personal computer and was arrested by FBI agents.[132]

- In April 2003, the United States Attorney's Office for the Northern District of California announced that Tse Thow Sun pled guilty on April 9, 2003, to theft of trade secrets. He admitted that in early 2002, while working for a language translation company, he delivered a laptop computer and a hard drive that contained trade secrets and confidential proprietary information to a competitor and asked for $3 million in payment. Mr. Sun, 32, a citizen of Singapore, was indicted by a federal Grand Jury on April 9, 2002. He was charged with theft of trade secrets, in violation of 18 U.S.C. §1832(a)(3); attempted theft of trade secrets, in violation of 18 U.S.C. §1832(a)(4); and interstate transportation of stolen goods, in violation of 18 U.S.C. §2314. Under the plea agreement, Mr. Sun pled guilty to theft of trade secrets.[133]
- In May 2003, three Swedish employees of LM Ericsson were charged with espionage for allegedly stealing intellectual property and sending it to Russian spies. "[Afshin] Bavand was arrested Nov. 5, 2002, while talking to a Russian intelligence agent in a Stockholm suburb. Police searched the Russian, who wasn't identified, and found $4,000 in cash and Ericsson documents."[134]
- The series of attacks codenamed *Titan Rain* was discovered by Shawn Carpenter in late 2003. Carpenter noticed a flood of expert hacker activity focusing on data theft from a wide range of "the country's most sensitive military bases, defense contractors and aerospace companies." Carpenter discovered that "the attacks emanated from just three Chinese routers that acted as the first connection point from a local network to the Internet." Carpenter worked with U.S. Army and FBI investigators to learn more about the attacks and the attackers. According to Thornburgh, various analysts judge that "Titan Rain is thought to rank among the most pervasive cyberespionage threats that U.S. computer networks have ever faced."[135]
- In July 2004, an Indian software engineer employed by a U.S. company's software development center in India was accused of "zipping up" proprietary software source code for printing identification cards and uploading it to her personal e-mail account. Jolly Technologies shut down its Mumbai operations as a result of the breach of security.[136]
- In 2005 and 2006, EMC filed lawsuits against several employees for allegedly stealing trade secrets.[137]
- In December 2006, two Chinese nationals, Fei Ye and Ming Zhong, pleaded guilty in December 2006 to charges of economic espionage on behalf of the People's Republic of China. They were arrested in November 2001 with stolen trade secrets in their luggage; the information was taken from Sun Microsystems and Transmeta Corporation. The agents were planning to design a competing microprocessor using the stolen designs; profits were to have been shared with the City of Hangzhou and the Province of Zhejiang. The agents' company was funded in part by the National High Technology Research and Development Program of China.[138]
- In April 2008, sleeper agent Chi Mak, a naturalized U.S. citizen who lived peacefully in Los Angeles for 20 years, was sentenced to 24.5 years in federal prison for industrial espionage. He stole detailed plans for U.S. Navy equipment including

submarine propulsion systems and tried to send them to China via his brother and sister-in-law.[139]

- In 2009, Siobhan Gorman, writing in *The Wall Street Journal*, reported as follows:

> Cyberspies have penetrated the U.S. electrical grid and left behind software programs that could be used to disrupt the system, according to current and former national-security officials. The spies came from China, Russia, and other countries, these officials said, and were believed to be on a mission to navigate the U.S. electrical system and its controls. The intruders haven't sought to damage the power grid or other key infrastructure, but officials warned they could try during a crisis or war. "The Chinese have attempted to map our infrastructure, such as the electrical grid," said a senior intelligence official. "So have the Russians." The espionage appeared pervasive across the U.S. and doesn't target a particular company or region, said a former Department of Homeland Security official. "There are intrusions, and they are growing," the former official said, referring to electrical systems. "There were a lot last year."[140]

- The Office of the National Counterintelligence Executive (ONCIX) published its Report to Congress on Foreign Economic Collection and Industrial Espionage 2009–2011 with the title "Foreign Spies Stealing U.S. Economic Secrets in Cyberspace." The Executive Summary included this commentary:

> Sensitive US economic information and technology are targeted by the intelligence services, private sector companies, academic and research institutions, and citizens of dozens of countries.

  - Chinese actors are the world's most active and persistent perpetrators of economic espionage. U.S. private-sector firms and cybersecurity specialists have reported an onslaught of computer network intrusions that have originated in China, but the IC cannot confirm who was responsible.
  - Russia's intelligence services are conducting a range of activities to collect economic information and technology from U.S. targets.
  - Some U.S. allies and partners use their broad access to U.S. institutions to acquire sensitive U.S. economic and technology information, primarily through aggressive elicitation and other human intelligence (HUMINT) tactics. Some of these states have advanced cyber capabilities.[141]
- A March 2012 report detailed how a successful supervisory control and data acquisition (SCADA) software company, American Superconductor Corporation (AMSC), was practically destroyed economically by its major customer, the Chinese Sinovel company, which stole its proprietary wind-turbine software and then stopped paying for any further software services.[142]
- By early 2013, Symantec's 2012 *Internet Security Threat Report*, Vol. 18 reported that small businesses were increasingly targeted for cyberattacks and industrial espionage: "In 2012, 50 percent of all targeted attacks were aimed at businesses with fewer than 2,500 employees. In fact, the largest growth area for targeted attacks in 2012 was businesses with fewer than 250 employees; 31 percent of all attacks targeted them."[143]

2.16 CONCLUDING REMARKS. At some point, history becomes current events. At the time of writing (May 2013), the trends we were seeing dimly when the fifth edition of this work was published have become clearer. As the second decade

of the 21st century reaches its midpoint, organized crime has become an integral part of the computer-crime scene—and vice versa. The Russian criminal underworld has increasingly invested in high-technology forms of fraud and also relies on high-tech communications for marketing of criminal undertakings, such as international traffic in drugs, armaments, and slaves. Information warfare has become a real issue as China advances in technology by stealing industrial secrets and capitalizing on the savings in research and development—and seeks growing global power. Terrorist groups cannot ignore the power of asymmetric warfare and must be presumed to be planning attacks on critical infrastructures worldwide. As the global communications network spreads throughout the world, governments, corporations, and individuals will have to increase their collaboration and vigilance to defeat the growing army of computer criminals of every type.

2.17 FURTHER READING

Anderson, N. *The Internet Police: How Crime Went Online, and the Cops Followed*. W. W. Norton & Company, 2013.

Banks, M. A. *Web Psychos, Stalkers and Pranksters: How to Protect Yourself in Cyberspace*. Coriolis Group Books, 1997.

Bequai, A. *Technocrimes: The Computerization of Crime and Terrorism*. Lexington Books, 1987.

Freedman, D. H., and C. C. Mann. *@Large: The strange case of the world's biggest Internet invasion*. Simon & Schuster, 1997.

Goodell, J. *The Cyberthief and the Samurai: The True Story of Kevin Mitnick—and the Man Who Hunted Him Down*. Dell, 1996.

Hafner, K., and J. Markoff. *Cyberpunk: Outlaws and Hackers on the Computer Frontier*. Simon & Schuster, 1991.

Hitchcock, J. A. *True Crime Online: Shocking Stories of Scamming, Stalking, Murder, and Mayhem*. Information Today, 2012.

Johnson, M. *Cyber Crime, Security and Digital Intelligence*. Gower Publishing, 2013.

Levy, S. *Hackers: Heroes of the Computer Revolution*. Doubleday, 1984.

Littman, J. *The Watchman: The Twisted Life and Crimes of Serial Hacker Kevin Poulson*. Little, Brown, 1997.

Menn, J. *Fatal System Error: The Hunt for the New Crime Lords Who Are Bringing Down the Internet*. PublicAffairs, 2010.

Mitnick, K. *Ghost in the Wires: My Adventures as the World's Most Wanted Hacker*. Little, Brown, 2011.

Mungo, P. *Approaching Zero: The Extraordinary Underworld of Hackers, Phreakers, Virus Writers, and Keyboard Criminals*. Random House, 1993.

Parker, D. B. *Fighting Computer Crime: A New Framework for Protecting Information*. John Wiley & Sons, 1998.

Poulson, K. *Kingpin: How One Hacker Took Over the Billion-Dollar Cybercrime Underground*. Crown, 2011.

Power, R. *Tangled Web: Tales of Digital Crime from the Shadows of Cyberspace*. Indianapolis: Que, 2000.

Shackelford, S. J. *Managing Cyber Attacks in International Law, Business, and Relations: In Search of Cyber Peace*. Cambridge University Press, 2013.

Shimomura, T., and J. Markoff. *Takedown: The Pursuit and Capture of Kevin Mitnick, America's Most Wanted Computer Outlaw—by the Man Who Did It*. Hyperion, 1996.

Slatalla, M., and J. Quittner. *Masters of Deception: The Gang that Ruled Cyberspace*. HarperCollins, 1995.

Sterling, B. *The Hacker Crackdown: Law and Disorder on the Electronic Frontier*. Bantam Doubleday Dell, 1992.

Stoll, C. *The Cuckoo's Egg: Tracking a Spy through the Maze of Computer Espionage*. Simon & Schuster, 1989.

Taylor, R. W., E. J. Fritsch, J. R. Liederbach, and T. J. Holt. *Digital Crime and Digital Terrorism*, 2nd ed. Prentice-Hall, 2010.

Webb, W. *You've Been Hacked: 15 Hackers You Hope Your Computer Never Meets*. CreateSpace Independent Publishing Platform, 2013.

Wells, J. T. *Computer Fraud Casebook: The Bytes that Bite.* Wiley, 2009.

2.18 NOTES

1. Some of the materials in this chapter use text from the author's prior publications, to which he holds the copyright. However, specific attributions or quotation marks in such cases are generally avoided, because changes are extensive and the typographical notations marking the changes would have been intrusive and disruptive.
2. Concordia University, "Who We Are: History," 2008, www.concordia.ca/about/whoweare/ourhistory/sgw.php
3. T. Whiteside, *Computer Capers: Tales of Electronic Thievery, Embezzlement, and Fraud* (New York: New American Library, 1978).
4. J. Gehl and S. Douglas, "Survey Reveals Epidemic of Battered PCs," *NewsScan*, June 5, 2001.
5. M. Delio, "Battered Computers: An Epidemic," *Wired*, June 5, 2001, www.wired.com/culture/lifestyle/news/2001/06/44284
6. NIPC/DHS, "Physical Attack Still the Biggest Threat," *Daily Open-Source Threat Report*, April 11, 2003.
7. T. Fricke, "Physical Security of Electronic Voting Terminals," *RISKS* 23, No. 20 (2004), http://catless.ncl.ac.uk/Risks/23.30.html
8. Whiteside, *Computer Capers*.
9. S. Spielberg, director, *Catch Me If You Can*, 2002, www.imdb.com/title/tt0264464/
10. R. Bell, *Skywayman: The Story of Frank W. Abagnale, Jr.* (Crime Library: Criminal Minds and Methods, 2008), www.trutv.com/library/crime/criminal_mind/scams/frank_abagnale/index.html
11. Whiteside, *Computer Capers*.
12. T. C. Greene, "Chapter One: Kevin Mitnick's Story." *The Register*, January 13, 2003, www.theregister.co.uk/2003/01/13/chapter_one_kevin_mitnicks_story/
13. J. Littman, *The Fugitive Game: Online with Kevin Mitnick—The Inside Story of the Great Cyberchase* (Boston: Little, Brown, 1996), p. 30.
14. A. N. Mayorkas and T. Mrozek, "Kevin Mitnick Sentenced to Nearly Four Years in Prison; Computer Hacker Ordered to Pay Restitution to Victim Companies Whose Systems Were Compromised," Press Release, U.S. Department of Justice, United States Attorney's Office, Central District of California, August 9, 1999, www.usdoj.gov/criminal/cybercrime/mitnick.htm

15. P. Jacobus, "Mitnick Released from Prison," *CNET News*, September 21, 2000, http://news.cnet.com/Mitnick-released-from-prison/2100-1023_3-235933.html
16. K. D. Mitnick and W. L. Simon, *The Art of Intrusion: The Real Stories Behind the Exploits of Hackers, Intruders & Deceivers* (New York: John Wiley & Sons, 1995). K. D. Mitnick and W. L. Simon, *The Art of Deception: Controlling the Human Element of Security* (Hoboken, NJ: John Wiley & Sons, 2003). J. Long, J. Wiles, and K. D. Mitnick, *No Tech Hacking: A Guide to Social Engineering, Dumpster Diving, and Shoulder Surfing* (Syngress, 2008).
17. E. Corley, director (as "Emmanuel Goldstein"), *Freedom Downtime* (2002), www.imdb.com/title/tt0309614/
18. R. Davies, "Origins of Money and of Banking," 2005, www.projects.ex.ac.uk/RDavies/arian/origins.html
19. "Origin and History of Credit Cards," Financial Web: Credit Cards, 2008, www.finweb.com/banking-credit/origin-and-history-of-credit-cards.html or http://tinyurl.com/5c2yhj
20. J. Rosenberg, "The First Credit Card," About.com: 20th-Century History, 2008, http://history1900s.about.com/od/1950s/a/firstcreditcard.htm or tinyurl.com/6en9kg
21. B. Hutchins, "Notes on the Fair Credit Billing Act (FCBA)," 2002, www.ftc.gov/os/comments/dncpapercomments/04/lsap7.pdf
22. L. S. Fox, ed., *The Federal Reserve System: Purposes & Functions*, 9th ed. (Washington, DC: Board of Governors of the Federal Reserve System, 2005), www.federalreserve.gov/pf/pdf/pf_1.pdf; Chapter 2, "Consumer and Community Affairs," p. 78 (p. 4 of PDF file). www.federalreserve.gov/pf/pdf/pf_6.pdf
23. "Origin and History of Credit Cards."
24. K. Shorter, "Plastic Payments: Trends in Credit Card Fraud," FBI Law Enforcement Bulletin (June 1997), www.fbi.gov/publications/leb/1997/june971.htm
25. www.usdoj.gov, news release of August 5, 2008.
26. Lynn Langton, "Identity Theft Reported by Households, 2005–2010," Bureau of Justice Statistics, November 30, 2011, www.bjs.gov/index.cfm?ty=pbdetail&iid=2207
27. Javelin Strategy & Research, "2013 IDENTITY FRAUD REPORT: Data Breaches Becoming a Treasure Trove for Fraudsters." Javelin, 2013, www.javelinstrategy.com/brochure/276
28. B. Sterling, *The Hacker Crackdown: Law and Disorder on the Electronic Frontier* (New York: Bantam, 1992). Available free online: www.mit.edu/hacker/hacker.html
29. E. McCracken, "Dial-Tone Phreak," *New York Times*, December 30, 2007, www.nytimes.com/2007/12/30/magazine/30joybubbles-t.html?ex=1356584400&en=8d26486125a53d83&ei=5124&partner=permalink&exprod=permalink or http://tinyurl.com/5s49cu
30. John T. Draper home page, www.webcrunchers.com
31. S. Wozniak and G. Smith, i*Woz: Computer Geek to Cult Icon: How I Invented the Personal Computer, Co-Founded Apple, and Had Fun Doing It* (New York: Norton, 2006).
32. R. Rosenbaum, "Secrets of the Little Blue Box," *Esquire Magazine* (October 1971). Reprinted at www.slate.com/articles/technology/the_spectator/2011/10/the_article_that_inspired_steve_jobs_secrets_of_the_little_blue_.html

33. The text displayed was available on Poulsen's Website until at least April 5, 2001, according to the Internet Archive. Sometime after that date, the biography was shortened and then sometime on or before December 4, 2002, it disappeared altogether and was replaced by a redirect to search for the string "By Kevin Poulsen" in Google.
34. K. Poulsen, "MySpace Predator Caught by Code," *Wired*, October 16, 2006, www.wired.com/science/discoveries/news/2006/10/71948
35. *Wired* magazine, "Threat Level" blog, www.wired.com/threatlevel
36. B. Trumbore, "Ray Dirks and the Equity Funding Scandal," *Wall Street History*, February 6, 2004, www.stocksandnews.com/wall-street-history.php?aid=MTU3M19XUw==
37. M. Kabay, "Crime, Use of Computers in." In H. Bidgoli, ed., *Encyclopedia of Information Systems*, vol. 1 (New York: Academic Press, 2003), www2.norwich.edu/mkabay/overviews/crime_use_of_computers_in.pdf or http://tinyurl.com/3wqfxc
38. D. E. Hoffman, "CIA Slipped Bugs to Soviets: Memoir Recounts Cold War Technological Sabotage," *Washington Post,* February 27, 2004, www.msnbc.msn.com/id/4394002
39. E. D. Shaw, K. G. Ruby, and J. M. Post, "The Insider Threat to Information Systems," *Security Awareness Bulletin* No. 2-98, Department of Defense Security Institute (September 1998), www.ntc.doe.gov/cita/CI_Awareness_Guide/Treason/Infosys.htm
40. CNN.com, Y2K Archive, "Looking at the Y2K Bug," 2000, www.cnn.com/TECH/specials/y2k/ (URL inactive)
41. Gardica-Feijóo, L., and J. R. Wingender, "Y2K: Myth or Reality." *Quarterly Journal of Business and Economics* (Summer 2007), http://findarticles.com/p/articles/mi_qa5466/is_200707/ai_n21295780/pg_1 or http://tinyurl.com/64w5jm (URL inactive)
42. United States Department of Homeland Security, "Look Before You Click: Trojan Horses and Other Attempts to Compromise Networks," *Joint Information Bulletin*, December 21, 2005, www.us-cert.gov/reading_room/JIB-Trojan122105.pdf or http://tinyurl.com/6zwmes (URL inactive)
43. D. Izenberg, "Trojan Horse Masterminds Being Extradited to Israel," *Jerusalem Post*, January 18, 2006. Available for purchase online: http://pqasb.pqarchiver.com/jpost/access/972012371.html?dids=972012371:972012371&FMT=ABS&FMTS=ABS:FT&type=current&date=Jan+18%2C+2006&author=DAN+IZENBERG&pub=Jerusalem+Post&edition=&startpage=04&desc=%27Trojan+horse%27+heads+extradited+to+Israel or http://tinyurl.com/5wlsgz (URL inactive)
44. W. K. Haskins, "Married Couple Indicted for Corporate Espionage," SCI-TECH TODAY.com, March 7, 2006, www.sci-tech-today.com/story.xhtml?story_id=12100DICT7FG&page=1 or http://tinyurl.com/3qantt (URL inactive)
45. L. Leyden, "Spyware-for-Hire Couple Plead Guilty: Israeli Prison Looms for Haephratis," *The Register*, March 15, 2006, www.theregister.co.uk/2006/03/15/spyware_trojan_guilty_plea/
46. "Court Hands Hefty Fine and Jail Sentence to Israeli Spyware Couple, Reports Sophos," *Sophos*, March 27, 2006, www.sophos.com/pressoffice/news/articles/2006/03/israelspyduo.html or http://tinyurl.com/4gx38p

47. J. Leyden, "Israeli Spyware-for-Hire PIs Jailed," *The Register*, April 29, 2008, www.theregister.co.uk/2008/04/29/spyware-for-hire/
48. R. Stiennon, "Four Private Investigators in the Israeli Trojan Fiasco Sentenced. Finally," *Network World*, "Stiennon on Security," April 30, 2008, www.networkworld.com/community/node/27387
49. J. L. Tkacik, "Trojan Dragon: China's Cyber Threat," Heritage Foundation Backgrounder #2106, February 8, 2008, www.heritage.org/Research/asiaandthepacific/bg2106.cfm
50. T. Claburn, "Operation 'Cisco Raider' Nets $76 in Fake Gear: The Multiyear Effort to Curb the Flow of Counterfeit Network Hardware into the U.S. and Canada Reflects a Steady Escalation in the War on Intellectual Property Crime," *InformationWeek*, February 29, 2008, www.informationweek.com/operation-cisco-raider-nets-76-million-i/206901053
51. J. Markoff, "Trojan Horse Threat Stalks Pentagon after Bogus Hardware Purchase," *CIO TODAY*, May 12, 2008, www.cio-today.com/story.xhtml?story_id=103006ROXFYH or http://tinyurl.com/5tvz32 (URL inactive)
52. "Ms Smith," "DHS: Imported Tech Tainted with Backdoor Attack Tools," *NetworkWorld* | Privacy and Security Fanatic, July 12, 2011, www.networkworld.com/community/blog/dhs-imported-tech-tainted-backdoor-attack-too
53. For a detailed and personal view of malware history, see virus expert Roger Thompson's "Malicious Code," Chapter 2 in S. Bosworth and M. E. Kabay, eds. *Computer Security Handbook*, 4th ed. (Hoboken, NJ: John Wiley & Sons, 2002). Also, see Chapter 16 in this *Handbook*.
54. R. H. Zakon, "Hobbes' Internet Timeline v8.2," 1996, www.zakon.org/robert/internet/timeline/
55. "Virus Encyclopedia: History of Malware." Viruslist.com, 2008, www.viruslist.com/en/viruses/encyclopedia?chapter=153310937
56. "Virus Encyclopedia: History of Malware."
57. Clausthal University of Technology homepage (English), www.tu-clausthal.de/Welcome.php.en (URL inactive)
58. Thompson, "Malicious Code."
59. R. Patterson, "Re: IBM Christmas Virus," *Risks Forum Digest* 5, No. 80 (December 21, 1987):1.1, catless.ncl.ac.uk/Risks/5.80.html~subj1.1
60. C. Schmidt and T. Darby "The What, Why and How of the 1988 Internet Worm," 1995, snowplow.org/tom/worm/worm.html
61. Robert Morris MIT faculty biography, www.csail.mit.edu/user/972
62. D. Emm, "Changing Threats, Changing Solutions: A History of Viruses and Antivirus," Viruslist.com (now Securelist.com), April 14, 2008, www.securelist.com/en/analysis?pubid=204791996
63. "The WildList Organization International: Frequently Asked Questions." WildList Organization International, 2008, www.wildlist.org/faq.htm
64. R. Thompson, personal communication, May 25, 2008.
65. M. E. Kabay, "INFOSEC Year in Review 1999," 1999, www2.norwich.edu/mkabay/iyir/1999.PDF (URL inactive)
66. CERT Advisory CA-2000-04 Love Letter Worm. CERT/CC, May 9, 2000, www.cert.org/advisories/CA-2000-04.html

67. D. I. Hopper, "Focus of 'ILOVEYOU' Investigation Turns to Owner of Apartment," CNN.com, May 10, 2000, http://archives.cnn.com/2000/TECH/computing/05/10/i.love.you.03/index.html or http://tinyurl.com/4elq2l
68. "Suspected Creator of 'ILOVEYOU' Virus Chats Online," CNN.com chat transcript, September 26, 2000, http://archives.cnn.com/2000/TECH/computing/09/26/guzman.chat/
69. M. Landler, "A Filipino Linked to 'Love Bug' Talks about His License to Hack," *New York Times*, October 21, 2000, http://query.nytimes.com/gst/fullpage.html?res=990DE5D8113EF932A15753C1A9669C8B63 or http://tinyurl.com/4b826p
70. R. G. Smith, "Impediments to the Successful Investigation of Transnational High Tech Crime," *Trends & Issues in Crime and Criminal Justice*, No. 285 (December 13, 2004), www.crime-research.org/articles/trends-and-issues-in-criminal-justice/ or http://tinyurl.com/44pn4s
71. Jaikumar Vijayan, "Stuxnet Renews Power Grid Security Concerns: First Known SCADA Malware Program to Target Control Systems Prompts New Questions about Security of U.S. Power Grid," *NetworkWorld*, July 26, 2010, www.networkworld.com/news/2010/072610-stuxnet-renews-power-grid-security.html
72. Jim Finkle, "Researchers say Stuxnet was deployed against Iran in 2007." Reuters, February 26, 2013, www.reuters.com/article/2013/02/26/us-cyberwar-stuxnet-idUSBRE91P0PP20130226
73. A. Lawrence, "Internet Growing Pains—The Canter & Siegel Story," *Computer Business Review* (June 1994), www.coin.org.uk/roadshow/presentation/canter.html
74. K. K. Campbell, "A NET.CONSPIRACY SO IMMENSE.... Chatting with Martha Siegel of the Internet's Infamous Canter & Siegel," 1994, http://lcs.www.media.mit.edu/people/foner/Essays/Civil-Liberties/Project/green-card-lawyers.html or http://tinyurl.com/45f3fe (URL inactive)
75. D. R. Hilton, "Green Card Lottery—Last Call," 1994, http://groups.google.com/group/misc.legal/msg/3416cd3d6cfcdebe (URL inactive)
76. L. Flynn, "'Spamming' on the Internet," *New York Times*, October 16, 1994, www.nytimes.com/1994/10/16/business/sound-bytes-spamming-on-the-internet.html
77. A. Craddock, "Spamming Lawyer Disbarred," *Wired*, July 10, 1997, www.wired.com/politics/law/news/1997/07/5060
78. N. Swidey, "Spambusters: Cyberwarriors of many stripes have joined the battle against junk email. But the enemy is wily, elusive—and multiplying," *Boston Globe*, October 5, 2003, www.boston.com/news/globe/magazine/articles/2003/10/05/spambusters?mode=PF or http://tinyurl.com/4y3chj
79. C. Garretson, "The Summer of Spam: Record Growth, Record Irritation," *Network World*, August 16, 2007, www.networkworld.com/news/2007/081607-spam-summer.html or http://tinyurl.com/6xoda3 (URL inactive)
80. J. Leyden, "Most Spam Comes from Just Six Botnets," *The Register*, February 29, 2008, www.theregister.co.uk/2008/02/29/botnet_spam_deluge/
81. See NetworkWorld's "Security Research Center" for up-to-date news about spam and phishing: www.networkworld.com/topics/security.html

82. T. Espiner, "Police Maintain Uneasy Relations with Cybervigilantes," *CNET News*, January 17, 2007, http://news.cnet.com/Police-maintain-uneasy-relations-with-cybervigilantes/2100-7348_3-6150817.html or http://tinyurl.com/6fjykr
83. "The Net's Most Wanted," *CNET News*, August 16, 1996, http://news.cnet.com/2100-1023-221580.html
84. L. Z. Koch, "Jacking in from the 'Spam in the Stocking' Port: Unamailer Delivers Christmas Grief," CyberWire Dispatch, December 26, 1996, www.petting-zoo.net/~deadbeef/archive/2122.html
85. Koch, "Jacking in from the 'Spam in the Stocking' Port."
86. "Unamailer Explains Bombings," *CNET News*, December 30, 1996, http://news.cnet.com/Unamailer-explains-bombings/2100-1017_3-258247.html or http://tinyurl.com/422kgc
87. M. Richtel and S. Robinson, "Several Web Sites Are Attacked on Day after Assault Shut Yahoo," *New York Times*, February 9, 2000, www.nytimes.com/library/tech/00/02/biztech/articles/09hack.html
88. E. Messmer, "Web Sites Unite to Fight Denial-of-Service War," *NetworkWorld*, September 25, 2000, www.networkworld.com/news/2000/0925userdefense.html?nf&_ref=858966935 or http://tinyurl.com/4cuvsf
89. "Today's FBI: Facts and Figures," 2003, www.fbi.gov/libref/factsfigure/factsfiguresapri2003.htm
90. See "The RMA Debate" for resources about "The Revolution in Military Affairs," www.comw.org/rma/fulltext/asymmetric.html (URL inactive)
91. For German-speakers or those with automated translation programs, see "FAQ—Über den Chaos Computer Club," May 27, 2004, www.ccc.de/de/faq
92. T. von Randow, "Bildschirmtext: A Blow against the System," translation from *Die Zeit*, November 30, 1984, www.textfiles.com/news/boh-20f8.txt
93. H. Nissenbaum, "Hackers and the Battle for Cyberspace," *Dissent* (Fall 2002), www.dissentmagazine.org/article/hackers-and-the-battle-for-cyberspace
94. Chaos Computer Club press release, "Chaos Computer Club Takes Legal Proceedings against the Voting Computer in Hesse," January 7, 2008, www.ccc.de/updates/2008/wahlcomputer-hessen
95. P. Elmer-Dewitt, "The 414 Gang Strikes Again: Pranksters disrupt a hospital, and nobody is laughing," *Time* (August 29, 1983), www.time.com/time/magazine/article/09171949797,00.html
96. P. Elmer-Dewitt, "Cracking Down: Hackers face tough new laws," *Time* (May 14, 1984), www.time.com/time/magazine/article/09171955290,00.html
97. At the time of writing (May 2008), the group's site (www.cultdeadcow.com) simply showed the words "BE BACK REAL SOON! / -xXx- cDc loves you with the fervor of a THOUSAND SUNS!! –xXx-" and a link to a YouTube video of a teenager playing a ukulele and singing. Consult the Internet Archive for historical snapshots of the site; web.archive.org/web/*/www.cultdeadcow.com/ (URL inactive)
98. S. Rat, "The infamous … GERBIL FEED BOMB: Striking fear into the hearts of model citizens everywhere …" cDc communications, 1985, http://web.archive.org/web/20050212092311/www.cultdeadcow.com/cDc_files/cDc-0001.html or http://tinyurl.com/44yyth

99. E. Messmer, "Bad Rap for Back Orifice 2000?" CNN.com, July 21, 1999, www.cnn.com/TECH/computing/9907/21/badrap.idg/
100. Sterling, *Hacker Crackdown*.
101. Hanna-Barbera, *Super Friends*, 1973, Internet Movie Database, www.imdb.com/title/tt0069641/
102. Hanna-Barbera, *Challenge of the Super Friends*, 1978, Internet Movie Database, www.imdb.com/title/tt0076994/
103. Sterling, *Hacker Crackdown*.
104. M. Slatalla and J. Quittner, "Gang War in Cyberspace," *Wired*, 2.12 (December 1994), www.wired.com/wired/archive/2.12/hacker.html
105. Datastream Cowboy, "MOD Indicted," *Phrack* 4, No. 40 (July 8, 1992): 13, www.phrack.com/issues.html?issue=40&id=13
106. J. Dibbell, "The Prisoner: Phiber Optik Goes Directly to Jail," *Village Voice* (January 12, 1994), www.juliandibbell.com/texts/phiber.html
107. J. Barone, "Manifesto." TechnoZen.com, 2000, www.technozen.com/manifesto.htm
108. The Mentor, "The Conscience of a Hacker," *Phrack* 1, No. 7(1986): 3; www.phrack.com/issues.html?issue=7&id=3#article
109. M. Slatalla and J. Quittner, *Masters of Deception: The Gang that Ruled Cyberspace* (New York: HarperCollins, 1995).
110. Sterling, *Hacker Crackdown*.
111. D. Charles, "'Innocent' Hackers Want Their Computers Back," *New Scientist*, No. 1820, May 9, 1992, p. 9; www.newscientist.com/article/mg13418201.400-innocent-hackers-want-their-computers-back-.html or http://tinyurl.com/3vw26e
112. Sterling, *Hacker Crackdown*.
113. S. Jackson, "SJ Games vs. the Secret Service," 2008, www.sjgames.com/SS/
114. D. Gans and K. Goffman, "Mitch Kapor & John Barlow Interview," Electronic Frontier Foundation, August 5, 1990, http://w2.eff.org/Misc/Publications/John_Perry_Barlow/HTML/barlow_and_kapor_in_wired_interview.html or http://tinyurl.com/4pgskr
115. S. Sparks, "Judge's Decision in Steve Jackson Games v. United States Secret Service," March 12, 1993, www.sjgames.com/SS/decision-text.html
116. D. Fisher, "The Long, Strange Trip of the L0pht," SearchSecurity.com, March 17, 2008, http://searchsecurity.techtarget.com/news/1305880/The-long-strange-trip-of-the-L0pht
117. M. Fitzgerald, "L0pht in Transition," CSO, April 17, 2007, www.csoonline.com/article/print/221192
118. Symantec News Release, "Symantec to Acquire @stake," September 16, 2004, www.symantec.com/press/2004/n040916b.html
119. United States Department of Justice, Computer Crime & Intellectual Property Section, "Shadowcrew Organization Called 'One-Stop Online Marketplace for Identity Theft': Nineteen Individuals Indicted in Internet 'Carding' Conspiracy," October 28, 2004, www.usdoj.gov/criminal/cybercrime/mantovaniIndict.htm (URL inactive)

120. United States Department of Justice, Computer Crime & Intellectual Property Section, "Six Defendants Plead Guilty in Internet Identity Theft and Credit Card Fraud Conspiracy," November 17, 2005, www.usdoj.gov/criminal/cybercrime/mantovaniPlea.htm (URL inactive)
121. G. A. White and R. W. Kern, "Cleveland, Ohio Man Sentenced to Prison for Bank Fraud and Conspiracy," U.S. Department of Justice, Eastern District of Pennsylvania, February 28, 2006, www.usdoj.gov/criminal/cybercrime/flurySent.htm
122. United States Attorney's Office, District of New Jersey, "'Shadowcrew' Identity Theft Ringleader Gets 32 Months in Prison," June 29, 2006, www.usdoj.gov/usao/nj/press/files/mant0629_r.htm (URL inactive)
123. Peter Warren, "Hunt for Russia's Web Criminals: The Russian Business Network—Which Some Blame for 60% of All Internet Crime—Appears To Have Gone To Ground. But, Asks Peter Warren, Has It Really Disappeared?" *The Guardian*, November 15, 2007, www.guardian.co.uk/technology/2007/nov/15/news.crime
124. Bizuel, David. "Russian Business Network study," bizuel.org, November 20, 2007, www.bizeul.org/files/RBN_study.pdf
125. David Goldman, "TheCyber Mafia Has Already Hacked You," CNNMoney, July 7, 2011, http://money.cnn.com/2011/07/27/technology/organized_cybercrime/index.htm
126. Anonymous, "Portal: Anonymous/Chanology," *Encyclopedia Dramatica*, October 29, 2012, https://encyclopediadramatica.se/Portal:Anonymous/Chanology
127. Gregory Ferenstein,"Anonymous Threatens Massive WikiLeaks-Style Exposure, Announced On Hacked Gov Site," *TechCrunch*, January 26, 2013. http://techcrunch.com/2013/01/26/anonymous-threatens-massive-wikileaks-style-exposure-announced-on-hacked-gov-site/
128. U.S. Department of Justice, "Eight Members of New York Cell of Cybercrime Organization Indicted in $45 million Cybercrime Campaign," U.S. Department of Justice | Eastern District of New York, May 9, 2013, www.justice.gov/usao/nye/pr/2013/2013may09.html
129. "Guilty Pleas in Trade Secret Case," *San Francisco Business Times*, April 27, 2001, www.bizjournals.com/eastbay/stories/2001/04/23/daily42.html
130. Andrew Backover, "Feds: Trio Stole Lucent's Trade Secrets," *USA Today*, May 3, 2001, http://usatoday30.usatoday.com/life/cyber/tech/2001-05-03-lucent-scientists-china.htm
131. "Trade-Secret Case Is Expanded," *New York Times*, April 12, 2002, www.nytimes.com/2002/04/12/technology/12LUCE.html
132. John Markoff, "Silicon Valley Concern Says It Thwarted Software Theft," *New York Times*, September 20, 2002, www.nytimes.com/2002/09/20/technology/20SOFT.html
133. U.S. Department of Justice, "Chicago, Illinois Man Pleads Guilty to Theft of Trade Secrets, Offered to Sell Online Interpreter's Information," U.S. Department of Justice | Northern District of California, April 11, 2003, www.justice.gov/criminal/cybercrime/press-releases/2003/sunPlea.htm
134. "Three Charged in Ericsson Spy Investigation in Sweden," *USA Today*, May 8, 2003, http://usatoday30.usatoday.com/tech/news/2003-05-08-ericsson_x.htm

135. Nathan Thornburgh, "Inside the Chinese Hack Attack," *Time*, August 25, 2005, www.time.com/time/nation/article/085991098371,00.html

136. John Ribeiro, "Source Code Stolen from U.S. Software Company in India: Jolly Technologies Blamed an Insider for the Theft," *Computerworld*, August 5, 2004, www.computerworld.com/s/article/95045/Source_code_stolen_from_U.S._software_company_in_India?taxonomyId=082

137. Todd Wallack, "EMC Sues Ex-Employees To Guard Trade Secrets," *Boston Business Journal*, October 9, 2006, www.bizjournals.com/boston/stories/2006/10/09/story4.html?page=all

138. U.S. Department of Justice, "Two Men Plead Guilty to Stealing Trade Secrets from Silicon Valley Companies to Benefit China," U.S. Department of Justice | Northern District of California, December 14, 2006, www.justice.gov/criminal/cybercrime/press-releases/2006/yePlea.htm

139. Joby Warrick and Carrie Johnson, "Chinese Spy 'Slept' In U.S. for 2 Decades: Espionage Network Said To Be Growing," *Washington Post*, April 3, 2008, www.washingtonpost.com/wp-dyn/content/story/2008/04/02/ST2008040204050.html

140. Siobhan Gorman, "Electricity Grid in U.S. Penetrated By Spies," *Wall Street Journal*, April 8, 2009, http://online.wsj.com/article/SB123914805204099085.html

141. Office of the National Counterintelligence Executive (ONCIX), "Foreign Spies Stealing US Economic Secrets in Cyberspace: Report to Congress on Foreign Economic Collection and Industrial Espionage, 2009–2011," Office of the National Counterintelligence Executive, November 3, 2011, www.ncix.gov/publications/reports/fecie_all/Foreign_Economic_Collection_2011.pdf

142. Michael Riley and Ashlee Vance, "Inside the Chinese Boom in Corporate Espionage," *Bloomberg Businessweek*, March 15, 2012, www.businessweek.com/articles/2012-03-14/inside-the-chinese-boom-in-corporate-espionage

143. Symantec, "Internet Security Threat Report 2013," Symantec, April 15, 2013, www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v18_2012_21291018.en-us.pdf

CHAPTER 3

TOWARD A NEW FRAMEWORK FOR INFORMATION SECURITY*

Donn B. Parker, CISSP

3.1 PROPOSAL FOR A NEW INFORMATION SECURITY FRAMEWORK 3·1

3.2 SIX ESSENTIAL SECURITY ELEMENTS 3·3

3.2.1 Loss Scenario 1: Availability 3·4
3.2.2 Loss Scenario 2: Utility 3·4
3.2.3 Loss Scenario 3: Integrity 3·5
3.2.4 Loss Scenario 4: Authenticity 3·5
3.2.5 Loss Scenario 5: Confidentiality 3·6
3.2.6 Loss Scenario 6: Possession 3·7
3.2.7 Conclusions about the Six Elements 3·8

3.3 WHAT THE DICTIONARIES SAY ABOUT THE WORDS WE USE 3·9

3.4 COMPREHENSIVE LISTS OF SOURCES AND ACTS CAUSING INFORMATION LOSSES 3·10

3.4.1 Complete List of Information Loss Acts 3·11
3.4.2 Examples of Acts and Suggested Controls 3·14
3.4.3 Physical Information and Systems Losses 3·17
3.4.4 Challenge of Complete Lists 3·18

3.5 FUNCTIONS OF INFORMATION SECURITY 3·19

3.6 SELECTING SAFEGUARDS USING A STANDARD OF DUE DILIGENCE 3·20

3.7 THREATS, ASSETS, VULNERABILITIES MODEL 3·20

3.8 CONCLUSION 3·20

3.9 FURTHER READING 3·23

3.1 PROPOSAL FOR A NEW INFORMATION SECURITY FRAMEWORK.

Information security, historically, has been limited by the lack of a comprehensive, complete, and analytically sound framework for analysis and improvement. The persistence of the classic triad of CIA (confidentiality, integrity, availability) is inadequate to describe what security practitioners include and implement when doing their jobs. We need a new information security framework that is complete, correct, and consistent

*This chapter is a revised excerpt from Donn B. Parker, *Fighting Computer Crime* (New York: John Wiley & Sons, 1998), Chapter 10, "A New Framework for Information Security," pp. 229–255.

to express, in practical language, the means for information owners to protect their information from any adversaries and vulnerabilities.

The current focus on computer systems security is attributable to the understandable tendency of computer technologists to protect what they know best—the computer and network systems rather than the application of those systems. With a technological hammer in hand, everything looks like a nail. The primary security challenge comes from people misusing or abusing information, and often—but not necessarily—using computers and networks. Yet the individuals who currently dominate the information security folk art are neither criminologists nor computer application specialists.

This chapter presents a comprehensive new information security framework that resolves the problems of the existing models. The chapter demonstrates the need for six security elements—availability, utility, integrity, authenticity, confidentiality, and possession—to replace incomplete CIA security (which does not even seem to include security for information that is not confidential) in the new security framework. This new framework is used to list all aspects of security at a basic level. The framework is also presented in another form, the *Threats, Assets, Vulnerabilities Model*, which includes detailed descriptors for each topic in the model. This model supports the new security framework, demonstrating its contribution to advance information security from its current technological stage, and as a folk art, into the basis for an engineering and business art in cyberspace.

The new security framework model incorporates six essential parts:

1. **Security elements** of information to be preserved are:
   - Availability
   - Utility
   - Integrity
   - Authenticity
   - Confidentiality
   - Possession
2. **Sources** of loss of these security elements of information:
   - Abusers and misusers
   - Accidental occurrences
   - Natural physical forces
3. **Acts** that cause loss:
   - Destruction
   - Interference with use
   - Use of false data
   - Modification or replacement
   - Misrepresentations or repudiation
   - Misuse or failure to use
   - Location
   - Disclosure
   - Observation
   - Copying

- Taking
- Endangerment

4. **Safeguard functions** to protect information from these acts:
   - Audit
   - Avoidance
   - Deterrence
   - Detection
   - Prevention
   - Mitigation
   - Transference
   - Investigation
   - Sanctions and rewards
   - Recovery
5. **Methods** of safeguard selection:
   - Use due diligence
   - Comply with regulations and standards
   - Enable business
   - Meet special needs
6. **Objectives** to be achieved by information security:
   - Avoid negligence
   - Meet requirements of laws and regulations
   - Engage in successful commerce
   - Engage in ethical conduct
   - Protect privacy
   - Minimize impact of security on performance
   - Advance an orderly and protected society

In summary, this model is based on the goal of meeting owners' needs to protect the desired *security elements* of their information from sources of loss that engage in harmful *acts* and events by applying *safeguard functions* that are selected by accepted *methods* to achieve desired *objectives.* The sections of the model are explained next. It is important to note that security risk, return on security investment (ROSI), and net present value (NPV) based on unknown future losses and enemies and their intentions are not identified in this model, since they are not measurable and, hence, not manageable.

3.2 SIX ESSENTIAL SECURITY ELEMENTS. Six security elements in the proposed framework model are essential to information security. If any one of them is omitted, information security is deficient in protecting information owners. Six scenarios of information losses, all derived from real cases, are used to demonstrate this contention. We show how each scenario involves violation of one, and only one, element of information security. Thus, if we omit that element from information security, we also remove that scenario from the concerns of information security, which would

be unacceptable. It is likely that information security professionals will agree that all of these scenarios fall well within the range of the abuse and misuse that we need to protect against.

3.2.1 Loss Scenario 1: Availability. A rejected contract programmer, intent on sabotage, removed the name of a data file from the file directories in a credit union's computer. Users of the computer and the data file no longer had the file available to them, because the computer operating system recognizes the existence of information available for users only if it is named in the file directories. The credit union was shut down for two weeks while another programmer was brought in to find and correct the problem so that the file would be available. The perpetrator was eventually convicted of computer crime.

Except for availability, the other elements of information security—utility, integrity, authenticity, confidentiality, and possession—do not address this loss, and their state does not change in the scenario. The owner of the computer (the credit union) retained possession of the data file. Only the availability of the information was lost, but it is a loss that clearly should have been prevented by information security. Thus, the preservation of availability must be accepted as a purpose of information security.

It is true that good security practice might have prevented the disgruntled programmer from having use of the credit union application system, and credit union management could have monitored his work more carefully. They should not have depended on the technical capabilities and knowledge of only one person, and they should have employed several controls to preserve or restore the availability of data files in the computer, such as by maintaining a backup directory with the names of erased files and pointers to their physical location. The loss might have been prevented, or minimized, through good backup practices, good usage controls for computers and specific data files, use of more than one name to identify and find a file, and the availability of utility programs to search for files by content or to mirror file storage. These safeguards would at least have made the attack more difficult and would have confronted the programmer with the malfeasance of his act.

The severity of availability loss can vary considerably. A perpetrator may destroy copies of a data file in a manner that eliminates any chance of recovery. In other situations, the data file may be partially usable, with recovery possible for a moderate cost, or the user may have inconvenienced or delayed use of the file for some period of time, followed by complete recovery.

3.2.2 Loss Scenario 2: Utility. In this case, an employee routinely encrypted the only copy of valuable information stored in his organization's computer and then accidentally erased the encryption key. The usefulness of the information was lost and could be restored only through difficult cryptanalysis.

Although this scenario can be described as a loss of availability or authenticity of the encryption key, the loss focuses on the usefulness of the information rather than on the key, since the only purpose of the key was to facilitate encryption. The information in this scenario is available, but in a form that is not useful. Its integrity, authenticity, and possession are unaffected, and its confidentiality, unfortunately, is greatly improved.

To preserve utility of information in this case, management should require mandatory backup copies of all critical information and should control the use of powerful protective mechanisms such as cryptography. Management should require security walk-through tests during application development to limit unusable forms of information. It should minimize the adverse effects of security on information use and should

control the types of activities that enable unauthorized persons to reduce the usefulness of information.

The loss of utility can vary in severity. The worst-case scenario would be the total loss of usefulness of the information, with no possibility of recovery. Less severe cases may range from a partially useful state with the potential for full restoration of usefulness at moderate cost.

3.2.3 Loss Scenario 3: Integrity. In this scenario, a software distributor purchased a copy (on DVD) of a program for a computer game from an obscure publisher. The distributor made copies of the DVD and removed the name of the publisher from the DVD copies. Then, without informing the publisher or paying any royalties, the distributor sold the DVD copies in a foreign country. Unfortunately, the success of the program sales was not deterred by the lack of an identified publisher on the DVD or in the product promotional materials.

Because the DVD copies of the game did not identify the publisher that created the program, the copies lacked integrity. ("Integrity" means a state of completeness, wholeness, and soundness, or adhering to a code of moral values.) However, the copies did not lack authenticity, since they contained the genuine game program and only lacked the identity of the publisher, which was not necessary for the successful use of the product. Information utility of the DVD was maintained, and confidentiality and availability were not at issue. Possession also was not at issue, since the distributor bought the original DVD. But copyright protection was violated as a consequence of the loss of integrity and unauthorized copying of the otherwise authentic program.

Several controls can be applied to prevent the loss of information integrity, including using and checking sequence numbers, checksums, and/or hash totals to ensure completeness and wholeness for a series of items. Other controls include performing manual and automatic text checks for required presence of records, subprograms, paragraphs, or titles, and testing to detect violations of specified controls.

The severity of information integrity loss also varies. Significant parts of the information can be missing or misordered (but still available), with no potential for recovery. Or missing or misordered information can be restored, with delay and at moderate cost. In the least severe cases, an owner can recover small amounts of misordered or mislocated information in a timely manner at low cost.

3.2.4 Loss Scenario 4: Authenticity. In a variation of the preceding scenario, another software distributor obtained the program (on DVD) for a computer game from an obscure publisher. The distributor changed the name of the publisher on the DVD and in title screens to that of a well-known publisher, then made copies of the DVD. Without informing either publisher, the distributor then proceeded to distribute the DVD copies in a foreign country. In this case, the identity of a popular publisher on the DVDs and in the promotional materials significantly added to the success of the product sales.

Because the distributor misrepresented the publisher of the game, the program did not conform to reality: It was not an authentic game from the well-known publisher. Availability and utility are not at issue in this case. The game had integrity because it identified a publisher and was complete and sound. (Certainly the distributor lacked *personal* integrity because his acts did not conform to ethical practice, but that is not the subject of the scenario.) The actual publisher did not lose possession of the game, even though copies were deceptively represented as having come from a different publisher.

And, although the distributor undoubtedly tried to keep his actions secret from both publishers, confidentiality of the content of the game was not at issue.

What if someone misrepresents your information by claiming that it is his? Violation of CIA does not include this act. A stockbroker in Florida cheated his investors in a Ponzi (pyramid sales) scheme. He stole $50 million by claiming that he used a super-secret computer program on his giant computer to make profits of 60 percent per day by arbitrage, a stock trading method in which the investor takes advantage of a small difference in prices of the same stock in different markets. He showed investors the mainframe computer at a Wall Street brokerage firm and falsely claimed that it and the information stored therein were his, thereby lending believability to his claims of successful trading.

This stockbroker's scheme was certainly a computer crime, but the CIA elements do not address it as such because its definition of integrity does not include misrepresentation of information. "Integrity" means only that information is whole or complete; it does not address the validity of information. Obviously, confidentiality and availability do not cover misrepresentation either. The best way to extend CIA to include misrepresentation is to use the more general term "authenticity." We can then assign the correct English meaning to the phrase "integrity of information": wholeness, completeness, and good condition. Dr. Peter Neumann at SRI International is correct when he says that information with integrity means that the information is what you expect it to be. This does not, however, necessarily mean that the information is valid (you may expect it to be invalid). "Authenticity" is the word that means conformance to reality.

A number of controls can be applied to ensure authenticity of information. These include confirming transactions, names, deliveries, and addresses; validating products; checking for out-of-range or incorrect information; and using digital signatures and watermarks to authenticate documents.

The severity of authenticity loss can take several forms, including lack of conformance to reality with no recovery possible; moderately false or deceptive information with delayed recovery at moderate cost; or factually correct information with only annoying discrepancies. If the CIA elements included authenticity, with misrepresentation of information as an important associated threat, Kevin Mitnick (the notorious criminal hacker who used deceit as his principal tool for penetrating security barriers) might have faced a far more difficult challenge in perpetrating his crimes. The computer industry might have understood the need to prove computer operating system updates and Web sites genuine, to avoid misrepresentation with fakes before their customers used those fakes in their computers.

3.2.5 Loss Scenario 5: Confidentiality. A thief deceptively obtained information from a bank's technical maintenance staff. He used a stolen key to open the maintenance door of an automated teller machine (ATM) and secretly inserted a radio transmitter that he purchased from a Radio Shack store. The radio received signals from the touch-screen display in the ATM that customers use to enter their personal identification numbers (PINs) and to receive account balance information. The radio device broadcast the information to the thief's radio receiver in his nearby car, which recorded the PINs and account balances on tape in a modified videocassette recorder. The thief used the information to loot the customers' accounts from other ATMs. The police and the Federal Bureau of Investigation caught the thief after elaborate detective and surveillance efforts. He was sentenced to 10 years in a federal prison.

The thief violated the secrecy of the customers' PINs and account balances, and he violated their privacy. Availability, utility, integrity, and authenticity were unaffected in

this violation of confidentiality. The customers' and the bank's exclusive possession of the PINs and account balance information was lost, but not possession per se, because they still held and owned the information. Therefore, this was primarily a case of lost confidentiality.

According to most security experts, confidentiality deals with disclosure, but confidentiality also can be lost by observation, whether that observation is voluntary or involuntary, and whether the information is disclosed or not disclosed. For example, if you leave sensitive information displayed on an unattended computer monitor screen, you have disclosed it and it may or may not lose its confidentiality. If you turn the monitor off, leaving a blank screen, you have not disclosed sensitive information, but if someone turns the monitor on and reads its contents without permission, then confidentiality is lost by observation. We must prevent both disclosure and observation in order to protect confidentiality.

Controls to maintain confidentiality include using cryptography, training employees to resist deceptive social engineering attacks intended to obtain their technical knowledge, and controlling the use of computers and computer devices. Good security also requires that the cost of resources for protection not exceed the value of what may be lost, especially with low incidence. For example, protecting against radio frequency emanations in ATMs (as in this scenario) is probably not advisable, considering the cost of shielding and the paucity of such high-tech attacks.

The severity of loss of confidentiality can vary. The worst-case scenario loss is when a party with the intent and ability to cause harm observes a victim's sensitive information. In this case, unrecoverable damage may result. But information also may be known to several moderately harmful parties, with a moderate loss effect, or be known to one harmless, unauthorized party with short-term recoverable effect.

3.2.6 Loss Scenario 6: Possession. A gang of burglars aided by a disgruntled, recently fired operations supervisor broke into a computer center and stole tapes and disks containing the company's master files. They also raided the backup facility and stole all backup copies of the files. They then held the materials for ransom in an extortion attempt against the company. The burglary resulted in the company's losing possession of all copies of the master files as well as the media on which they were stored. The company was unable to continue business operations. The police eventually captured the extortionists with help from the company during the ransom payment, and they recovered the stolen materials. The burglars were convicted and served long prison sentences.

Loss of possession occurred in this case. The perpetrators delayed availability, but the company could have retrieved the files at any time by paying the ransom. Alternatively, the company could have re-created the master files from paper documents, but at great cost. Utility, integrity, and authenticity were not issues in this situation. Confidentiality was not violated because the burglars had no reason to read or disclose the files. Loss of ownership and permanent loss of possession would have been accomplished if the perpetrators had never returned the materials or if the company had stopped trying to recover them.

The security model must include protecting the possession of information so as to prevent theft, whether the information is confidential or not. Confidentiality, by definition, deals only with secret information that people may possess. Our increasing use of computers magnifies this difference; huge amounts of information are possessed for automated use and not necessarily held confidentially for only specified people to know. Computer object programs are examples of proprietary but not confidential

information we do not know but possess by selling, buying, bartering, giving, receiving, and trading until we ultimately control, transport, and use them. We have incorrectly defined possession if we include only the protective efforts for confidential material.

We protect the possession of information by preventing people from unauthorized taking, from making copies, and from holding or controlling it—whether confidentiality is involved or not. The loss of possession of information also includes the loss of control of it, and may allow the new possessor to violate its confidentiality at will. Thus, loss of confidentiality may accompany loss of possession. But we must treat confidentiality and possession separately to determine what actions criminals might take and what controls we need to apply to prevent their actions. Otherwise, we may overlook a particular threat or an effective control. The failure to anticipate a threat and vulnerability is one of the greatest dangers we face in security.

Controls that can protect the possession of information include using copyright laws, implementing physical and logical usage limitations, preserving and examining computer audit logs for evidence of stealing, inventorying tangible and intangible assets, using distinctive colors and labels on media containers, and assigning ownership to enforce accountability of organizational information assets.

The severity of loss of possession varies with the nature of the offense. In a worst-case scenario, a criminal may take information, as well as all copies of it, and there may be no means of recovery—either from the perpetrator or from other sources such as paper documentation. In a less harmful scenario, a criminal might take information for some period of time but leave some opportunity for recovery at a moderate cost. In the least harmful situation, an owner could possess more than one copy of information, leaving open the possibility of recovery from other sources (e.g., backup files) within a reasonable period of time.

3.2.7 Conclusions about the Six Elements. We need to understand some important differences between integrity and authenticity. For one, integrity deals with the intrinsic condition of information, while authenticity deals with the extrinsic value or meaning relative to external sources and uses. Integrity does not deal with the meaning of the information with respect to external sources, that is, whether the information is timely and not obsolete. Authenticity, in contrast, concerns the question of whether information is genuine or valid and not out of date with respect to its potential use. A user who enters false information into a computer possibly has violated authenticity, but as long as the information remains unchanged, it has integrity. An information security technologist who designs security into computer operating systems is concerned only with application information integrity because the designer cannot know if any user is entering false information. In this case, the security technologist's job is to ensure that both true and false information remain whole and complete. It is the information owner, with guidance from an information security advisor, who has the responsibility of ensuring that the information conforms to reality—in other words, that it has authenticity.

Some types of loss that information security must address require the use of all six elements of the framework model to determine the appropriate security to apply. Each of the six elements can be violated independently of the others, with one important exception: A violation of confidentiality always results in loss of exclusive possession, at the least. Loss of possession, however—even exclusive possession—does not necessarily result in loss of confidentiality.

Other than that exception, the six elements are unique and independent, and often require different security controls. Maintaining the availability of information does

not necessarily maintain its utility; information may be available but useless for its intended purpose, and vice versa. Maintaining the integrity of information does not necessarily mean that the information is valid, only that it remains the same or, at least, whole and complete. Information can be invalid and, therefore, without authenticity, yet it may be present and identical to the original version and, thus, have integrity. Finally, who is allowed to view and know information and who possesses it are often two very different matters.

Unfortunately, the written information security policies of many organizations do not acknowledge the need to address many kinds of information loss. This is because their policies are limited to achieving CIA. To define information security completely, the policies must address all six elements presented. Moreover, to eliminate (or at least reduce) security threats adequately, all six elements need to be considered to ensure that nothing is overlooked in applying appropriate controls. These elements are also useful for identifying and anticipating the types of abusive actions that adversaries may take—before such actions are undertaken.

For simplification and ease of reference, we can pair the six elements into three double elements, which should be used to identify threats and select proper controls, and we can associate them with synonyms so as to facilitate recall and understanding:

availability and utility	→	usability and usefulness
integrity and authenticity	→	completeness and validity
confidentiality and possession	→	secrecy and control

Availability and utility fit together as the first double element. Controls common to these elements include secure location, appropriate form for secure use, and usability of backup copies. Integrity and authenticity also fit together; one is concerned with internal structure and the other with conformance to external facts or reality. Controls for both include double entry, reasonableness checks, use of sequence numbers and checksums or hash totals, and comparison testing. Control of change applies to both as well. Finally, confidentiality and possession go together because, as discussed, they are interrelated. Commonly applied controls for both include copyright protection, cryptography, digital signatures, escrow, and secure storage.

The order of the elements here is logical, since availability and utility are necessary for integrity and authenticity to have value, and these first four elements are necessary for confidentiality and possession to have material meaning.

3.3 WHAT THE DICTIONARIES SAY ABOUT THE WORDS WE USE. CIA would be adequate for security purposes if the violation of confidentiality were defined to be anything done *with* information, if integrity were defined to be anything done *to* information, and if availability were to include utility, but these definitions would be incorrect and are not understood by many people. Information professionals are already defining the term "integrity" incorrectly, and we would not want to make matters worse. These definitions of security and the elements are relevant abstractions from *Webster's Third New International Dictionary* and *Webster's Collegiate Dictionary,* 10th edition.

Security—freedom from danger, fear, anxiety, care, uncertainty, doubt; basis for confidence; measures taken to ensure against surprise attack, espionage, observation, sabotage; resistance of a cryptogram to cryptanalysis usually measured by the time and effort needed to solve it.

Availability—present or ready for immediate use.

Utility—useful, fitness for some purpose.

Integrity—unimpaired or unmarred condition; soundness; entire correspondence with an original condition; adherence to a code of moral, artistic or other values; the quality or state of being complete or undivided; material wholeness.

Authenticity—quality of being authoritative, valid, true, real, genuine, worthy of acceptance or belief by reason of conformity to fact and reality.

Confidentiality—quality or state of being private or secret; known only to a limited few, containing information whose unauthorized disclosure could be prejudicial to the national interest.

Possession—act or condition of having or taking into one's control or holding at one's disposal; actual physical control of property by one who holds for himself, as distinguished from custody; something owned or controlled.

We lose credibility and confuse information owners if we do not use words precisely and consistently. When defined correctly, the six words are independent (with the exception that information possession is always violated when confidentiality is violated). They are also consistent, comprehensive, and complete. In other words, the six elements themselves possess integrity and authenticity, and therefore they have great utility. This does not mean that we will not find new elements or replace some of them as our insights develop and technology advances. (I first presented this demonstration of the need for the six elements in 1991 at the 14th U.S. National Security Agency/National Institute of Standards and Technology National Computer Security Conference in Baltimore.)

My definitions of the six elements are considerably shorter and simpler than the dictionary definitions, but appropriate for information security.

Availability—usability of information for a purpose

Utility—usefulness of information for a purpose

Integrity—completeness, wholeness, and readability of information and quality being unchanged from a previous state

Authenticity—validity, conformance, and genuineness of information

Confidentiality—limited observation and disclosure of knowledge

Possession—holding, controlling, and having the ability to use information

3.4 COMPREHENSIVE LISTS OF SOURCES AND ACTS CAUSING INFORMATION LOSSES.

The losses that we are concerned about in information security come from people who engage in unauthorized and harmful acts against information, communications, and systems, such as embezzlers, fraudsters, thieves, saboteurs, and criminal hackers. They engage in harmful using, taking, misrepresenting, observing, and every other conceivable form of human misbehavior. Natural physical forces such as air and earth movements, heat and cold, electromagnetic energy, living organisms, gravity and projectiles, and water and gases also are threats to information, as are inadvertent human errors.

Extensive lists of losses found in information security often include fraud, theft, sabotage, and espionage, along with disclosure, usage, repudiation, and copying. The first four losses in this list are criminal justice terms at a different level of abstraction from the last four and require an understanding of criminal law, which many information owners and security specialists lack. For example, fraud includes theft only if it is

performed using deception, and larceny includes burglary and theft from a victim's premises. What constitutes "premises" in an electronic network environment? This is a legal issue.

Many important types of information-related acts, such as false data entry, failure to perform, replacement, deception, misrepresentation, prolongation of use, delay of use, and even the obvious taking copies of information, are frequently omitted from lists of adverse incidents. Each of these losses may require different prevention and detection controls. This may be easily overlooked if our list of potential acts is incomplete—even though the acts that we typically omit are among the most common reported in actual loss experience. The people who cause losses often are aware that information owners have not provided adequate security and have not considered the full array of possible acts. It is, therefore, essential to include all types of potential harmful acts in our lists, especially when unique safeguards are applicable. Otherwise, we are in danger of being negligent, and those to whom we are accountable will view information security as incomplete or poorly conceived and implemented when a loss does occur.

The complete list of information loss acts in the next section is a comprehensive, nonlegalistic list of potential acts resulting in losses to or with information that I compiled from my 35 years in research about computer crime and security. I have simplified it to a single, low level of abstraction to facilitate understanding by information owners and to enable them to select effective controls. The list makes no distinction among the causes of the losses; as such, it applies equally well to accidental and intentional acts. Cause is largely irrelevant at this level of security analysis, as is the underlying intent or lack thereof. (Identifying cause is important at another level of security analysis. We need to determine the sources and motivation of threats in order to identify appropriate avoidance, deterrence, correction, and recovery controls.) In addition, the list makes no distinction between electronic and physical causes of loss, or among spoken, printed, or electronically recorded information.

The acts in the list are grouped to correspond to the six elements of information security outlined previously (e.g., availability and utility, etc.). Some types of acts in one element grouping may have a related effect in another grouping as well. For example, if no other copies of information exist, destroying the information (under *availability*) also may cause loss of possession, and taking (under *possession*) may cause loss of availability. Yet loss of possession and loss of availability are quite different, and may require different controls. I have placed acts in the most obvious categories, where a loss prevention analyst is likely to look first.

Here is an abbreviated version of the complete loss list for convenient use in the information security framework model:

- Destroy
- Interfere with use
- Introduce false data
- Modify or replace
- Misrepresent or repudiate

3.4.1 Complete List of Information Loss Acts

Availability and Utility Losses

- Destruction, damage, or contamination
- Denial, prolongation, acceleration, or delay in use or acquisition

- Movement or misplacement
- Conversion or obscuration

Integrity and Authenticity Losses

- Insertion, use, or production of false or unacceptable data
- Modification, replacement, removal, appending, aggregating, separating, or reordering
- Misrepresentation
- Repudiation (rejecting as untrue)
- Misuse or failure to use as required

Confidentiality and Possession Losses

- Locating
- Disclosing
- Observing, monitoring, and acquiring
- Copying
- Taking or controlling
- Claiming ownership or custodianship
- Inferring
- Exposing to all of the other losses
- Endangering by exposing to any of the other losses
- Failure to engage in or to allow any of the other losses to occur when instructed to do so

Users may be unfamiliar with some of the words in the lists of acts, at least in the context of security. For example, "repudiation" is a word that we seldom hear or use outside of the legal or security context. According to dictionaries, it means to refuse to accept acts or information as true, just, or of rightful authority or obligation. Information security technologists became interested in repudiation when the Massachusetts Institute of Technology (MIT) developed a secure network operating system for its internal use. The system was named Kerberos, taking the name of the three-headed dog that guarded the underworld in Greek mythology. Kerberos provides a means of forming secure links and paths between users and the computers serving them. Unfortunately, however, in early versions it allowed users to falsely deny using the links. This did not present any particular problems in the academic environment, but it did make Kerberos inadequate for business, even though its other security aspects were attractive. As the use of Kerberos spread into business, repudiation became an issue, and nonrepudiation controls became important.

Repudiation is an important issue in electronic transactions such as in electronic banking, purchases, and auctions used by so many people to automate their purchasing functions and Internet commerce, which require digital signatures, escrow, time stamps, and other authentication controls. I could, for example, falsely claim that I never ordered merchandise and that the order form or electronically transmitted ordering information that the merchant possesses is false. Repudiation is also a growing problem because of the difficulty of proving authorship or the source of electronic missives. And the inverse of repudiation—claiming that an act that did not happen actually did happen, or claiming that false information is true—is also important to security, although it

is often overlooked. Repudiation and its inverse are both types of misrepresentation, but I include both "repudiation" and "misrepresentation" on the list because they may require different types of controls.

Other words in the list of acts may seem somewhat obscure. For example, we seldom think of prolonging or delaying use as a loss of availability or a denial of use, yet they are losses that are often inflicted by computer virus attacks.

I use the word "locate" in the list rather than "access" because access can be confusing with regard to information security. Although it is commonly used in computer terminology, its use frequently causes confusion, as it did in the trial of Robert T. Morris for releasing the Internet worm of November 2, 1988, and in computer crime laws. For example, access may mean just knocking on a door or opening the door but not going in. How far "into" a computer must you go to "access" it? A perpetrator can cause a loss simply by locating information, because the owner may not want to divulge possession of such information. In this case, no access is involved. For these reasons, I prefer to use the terms "entry," "intrusion," and "usage"—as well as "locate"—to refer to a computer as the object of the action. I have a similar problem with the use of the word "disclosure" and ignoring observation as I indicated earlier. "Disclose" is a verb that means to divulge, reveal, make known, or report knowledge to others. We can disclose knowledge by:

- Broadcasting
- Speaking
- Displaying
- Showing
- Leaving it in the presence and view of another person
- Leaving it in possible view where another person is likely to be
- Handing or sending it to another person

Disclosure is what an owner or potential victim might do inadvertently or intentionally, not what a perpetrator does, unless it is the second act after stealing, such as selling stolen intellectual property to another person. Disclosure can be an abuse if a person authorized to know information discloses it to an unauthorized person, or if an unauthorized person discloses knowledge to another person without permission. In any case, confidentiality is lost or is potentially lost, and the person disclosing the information may be accused of negligence, violation of privacy, conspiracy, or espionage.

Loss of confidentiality also can occur by observation, whether the victim or owner disclosed knowledge, resisted disclosure, or did nothing either to protect or to disclose it. Observing is an abuse of listening, spying by eavesdropping, shoulder surfing (looking over another person's shoulder or overhearing), looking at or listening to a stolen copy of information, or even by tactile feeling, as in the case of reading Braille. We should think about loss of confidentiality as a loss caused by inadvertent disclosure by the victim, observation by the perpetrator, and disclosure by the perpetrator who passes information to a third party. Disclosure and observation of information that is not knowledge converts it into knowledge if cognition takes place. Disclosure always results in loss of confidentiality by putting information into a state where there is no longer any secrecy, but observation results in loss of confidentiality only if cognition or use to the detriment of the owner takes place. Privacy is a right that is a whole other topic that I do not cover here. (This issue is discussed in Chapter 69.)

Loss of possession of information (including knowledge) is the loss from the unintended or regretful giving or taking of information. At a higher level of crime description, we call it larceny (theft or burglary) or fraud (when deceit is involved). Possession seems to be most closely associated with confidentiality. The two are placed together in the list because they share the common losses of taking and copying (loss of exclusive possession). I could have used "ownership" of information, since it is a synonym for possession, but "ownership" seems to be not as broad, because someone may rightly or wrongly possess information that is rightfully owned by another. The concepts of owner or possessor of information, along with user, provider, or custodian of information, are important distinctions in security for assigning asset accountability. This provides another reason for including possession in the list.

The act of *endangerment* is quite different from, but applies to, the other losses. It means putting information in harm's way, or that a person has been remiss (and possibly negligent) by not applying sufficient protection to information, such as leaving sensitive or valuable documents in an unlocked office or open trash bin. Leaving a computer unnecessarily connected to the Internet is another example. Endangerment of information may lead to charges of negligence or criminal negligence and civil liability suits that may be more costly than direct loss incidents. My objectives of security in the framework model invoke a standard of due diligence to deal with this exposure.

The last act in the list—failure to engage in or allow any of the other acts when instructed to do so—may seem odd at first glance. It means that an information owner may require an act resulting in any of the other acts to be carried out. Or the owner may wish that an act be allowed to occur, or information to be put into danger of loss. There are occasions when information should be put in harm's way for testing purposes or to accomplish a greater good. For example, computer programmers and auditors often create information files that are purposely invalid for use as input to a computer to make sure that the controls to detect or mitigate a loss are working correctly. A programmer bent on crime might remove invalid data in a test input file to avoid testing a control that the perpetrator has neutralized or has avoided implementing for nefarious purposes. The list would surely be incomplete without this type of loss, yet I have never seen it included or discussed in any other information security text.

The acts in the list are described at the appropriate level for deriving and identifying appropriate security controls. At the next lower level of abstraction (e.g., read, write, and execute), the losses would not be so obvious and would not necessarily suggest important controls. At the level that I choose, there is no attempt to differentiate acts that make no change to information from those that do, since these differences are not important for identifying directly applicable controls or for performing threat analyses. For example, an act of modification changes the information, while an act of observation does not, but encryption is likely to be employed as a powerful primary control against both acts.

3.4.2 Examples of Acts and Suggested Controls.

The next examples illustrate the relationships between acts and controls in threat analysis. Groups of acts are followed by examples of the losses and applicable controls.

3.4.2.1 Destroy, Damage, or Contaminate.

Perpetrators or harmful forces can damage, destroy, or contaminate information by electronically erasing it, writing other data over it, applying high-energy radio waves to damage delicate electronic circuits, or physically damaging the media (e.g., paper, flash memory, or disks) containing it.

Controls include disaster prevention safeguards such as locked facilities, safe storage of backup copies, and write-usage authorization requirements.

3.4.2.2 Deny, Prolong, Delay Use or Acquisition. Perpetrators can make information unavailable by hiding it or denying its use through encryption and not revealing the means to restore it, or by keeping critical processing units busy with other work, such as in a denial-of-service attack. Such actions would not necessarily destroy the information. Similarly, a perpetrator may prolong information use by making program changes that slow the processing in a computer or by slowing the display of the information on a screen. Such actions might cause unacceptable timing for effective use of the information. Information acquisition may be delayed by requiring too many passwords to retrieve it or by slowing retrieval. These actions can make the information obsolete by the time it becomes available.

Controls include making multiple copies available from different sources, preventing overload of processing by selective allowance of input, or preventing the activation of harmful mechanisms such as computer viruses by using antiviral utilities.

3.4.2.3 Enter, Use, or Produce False Data. Data diddling, my term for false data entry and use, is a common form of computer crime, accounting for much of the financial and inventory fraud. Losses may be either intentional, such as those resulting from the use of Trojan horses (including computer viruses), or unintentional, such as those from input errors.

Most internal controls such as range checks, audit trails, separation of duties, duplicate data entry detection, program proving, and hash totals for data items protect against these threats.

3.4.2.4 Modify, Replace, or Reorder. These acts are often intelligent changes rather than damage or destruction. *Reordering,* which is actually a form of modification, is included separately because it may require specific controls that could otherwise be overlooked. Similarly, *replacement* is included because users might not otherwise include the idea of replacing an entire data file when considering modification. Any of these actions can produce a loss inherent in the threats of entering and modifying information, but including all of them covers modifying data both before entry and after entry, since each requires different controls.

Cryptography, digital signatures, usage authorization, and message sequencing are examples of controls to protect against these acts, as are detection controls to identify anomalies.

3.4.2.5 Misrepresent. The claim that information is something different from what it really is or has a different meaning from what was intended arises in counterfeiting, forgery, fraud, impersonation (of authorized users), and many other deceptive activities. Hackers use misrepresentation in social engineering to deceive people into revealing information needed to attack systems. Misrepresenting old data as new information is another act of this type.

Controls include user and document authentication methods such as passwords, digital signatures, and data validity tests. Making trusted people more resistant to deception by reminders and training is another control.

3.4.2.6 Repudiate. This type of loss, in which perpetrators generally deny having made transactions, is prevalent in electronic data interchange (EDI) and Internet

commerce. Oliver North's denial of the content of his email messages is a notable example of repudiation, but as I mentioned earlier, the inverse of repudiation also represents a potential loss.

Repudiation can be controlled most effectively through the use of digital signatures and public key cryptography. Trusted third parties, such as certificate authorities with secure computer servers, provide the independence of notary publics to resist denial of truthful information as long as they can be held liable for their failures.

3.4.2.7 Misuse or Fail to Use as Required. Misuse of information is clearly an act resulting in many information losses. Misuse by failure to perform duties such as updating files or backing up information is not so obvious and needs explicit identification. Implicit misuse by conforming exactly to inadequate or incorrect instructions is a sure way to sabotage systems.

Information usage control and internal application controls that constrain the modification or use of trusted software help to avoid these problems. Keeping secure logs of routine activities can help catch operational vulnerabilities.

3.4.2.8 Locate. Unauthorized use of someone's computer or data network to locate and identify information is a crime under most computer crime statutes—even if there is no overt intention to cause harm. Such usage is a violation of privacy, and trespass to engage in such usage is a crime under other laws.

Log-on and usage controls are major features in many operating systems such as Microsoft Windows and some versions of UNIX as well as in add-on security utilities such as *RACF* and *ACF2* for large IBM computers and many security products for personal computers.

3.4.2.9 Disclose. Preventing information from being revealed to people not authorized to know it is the purpose of business, personal, and government secrecy. Disclosure may be verbal, by mail, or by transferring messages or files electronically or on disks, flash memories, or tape. Disclosure can result in loss of privacy and trade secrets.

Military organizations have advanced protection of information confidentiality to an elaborate art form.

3.4.2.10 Observe or Monitor. Observation, which requires action on the part of a perpetrator, is the inverse of disclosure, which results from actions of a possessor. Workstation display screens, communication lines, and monitoring devices such as recorders and audit logs are common targets of observation and monitoring. Observation of output from printers is another possible source, as is shoulder surfing—the technique of watching screens of other computer users.

Physical entry protection for input and output devices represents the major control to prevent this type of loss. Preventing wiretapping and eavesdropping is also important.

3.4.2.11 Copy. Copy machines and the software *copy* command are the major sources of unauthorized copying. Copying is used to violate exclusive possession and privacy. Copying can destroy authenticity, as when used to counterfeit money or other business instruments.

Location and use controls are effective against copying, as are unique markings such as those used on U.S. currency and watermarks on paper and in computer files.

3.4.2.12 Take. Transferring data files in computers or networks constitutes taking. So does taking small computers and DVDs or documents for the value of the information stored in them. Perpetrators can easily take copies of information without depriving the owner of possession or confidentiality.

A wide range of physical and logical location controls applies to these losses; most are based on common sense and a reasonable level of due care.

3.4.2.13 Endanger. Putting information into locations or conditions in which others may cause loss in any of the previously described ways clearly endangers the information, and the perpetrator may be accused of negligence, at least.

Physical and logical means of preventing information from being placed in danger are important. Training people to be careful, and holding them accountable for protecting information, may be the most effective means of preventing endangerment.

3.4.3 Physical Information and Systems Losses. Information also can suffer from physical losses such as those caused by floods, earthquakes, radiation, and fires. Although these losses may not directly affect the information itself (e.g., knowledge of operating procedures held in the minds of operators), they can damage or destroy the media and the environment that contain representations of the information. Water, for example, can destroy printed pages and damage magnetic disks; physical shaking or radio frequency radiation can short-out electronic circuits, and fires can destroy all types of media. Overall, physical loss may occur in seven natural ways by application of:

1. Extreme temperature
2. Gases
3. Liquids
4. Living organisms
5. Projectiles
6. Movements
7. Energy anomalies

Each way, of course, comes from specific sources of loss (e.g., smoke or water). And the various ways can be broken down further, to identify the underlying cause of the source of loss. For example, the liquid that destroys information may be water flowing from a plumbing break above the computer workstation, caused in turn by freezing weather. The next list presents examples of each of the seven major sources of physical loss.

1. **Extreme temperature.** Heat or cold. Examples: sunlight, fire, freezing, hot weather, and the breakdown of air-conditioning equipment.
2. **Gases.** War gases, commercial vapors, humid or dry air, suspended particles. Examples: sarin nerve gas, PCBs from exploding transformers, release of Freon from air conditioners, smoke and smog, cleaning fluid, and fuel vapors.
3. **Liquids**. Water, chemicals. Examples: floods, plumbing failures, precipitation, fuel leaks, spilled drinks, acid and base chemicals used for cleaning, and computer printer fluids.

4. **Living organisms**. Viruses, bacteria, fungi, plants, animals, and human beings. Examples: sickness of key workers, molds, contamination from skin oils and hair, contamination and electrical shorting from defecation and release of body fluids, consumption of information media such as paper or of cable insulation, and shorting of microcircuits from cobwebs.
5. **Projectiles**. Tangible objects in motion, powered objects. Examples: meteorites, falling objects, cars and trucks, airplanes, bullets and rockets, explosions, and windborne objects.
6. **Movement**. Collapse, shearing, shaking, vibration, liquefaction, flows, waves, separation, slides. Examples: dropping or shaking fragile equipment, earthquakes, earth slides, lava flows, sea waves, and adhesive failures.
7. **Energy anomalies**. Electric surge or failure, magnetism, static electricity, aging circuitry; radiation, sound, light, radio, microwave, electromagnetic, atomic. Examples: electric utility failures, proximity of magnets and electromagnets, carpet static, electromagnetic pulses (EMP) from nuclear explosions, lasers, loudspeakers, high-energy radio frequency (HERF) guns, radar systems, cosmic radiation, and explosions.

Although meteorites, for example, clearly pose little danger to computers, it is nonetheless important to include all such unlikely events in a thorough analysis of potential threats. In general, include every possible act included in a threat analysis. Then consider it carefully; if it is too unlikely, document the consideration and discard the item. It is better to have thought of a source of loss and to have discarded it than to have overlooked an important one. Invariably, when you present a threat analysis to others, someone will try to surprise the developer with another source of loss that has been overlooked.

Insensitive practitioners have ingrained inadequate loss lists in the body of knowledge from the very inception of information security. Proposing a major change at this late date is a bold action that may take significant time to accomplish. However, we must not perpetuate our past inadequacies by using the currently accepted destruction, disclosure, use, and modification (DDUM) as a complete list of losses. We must not underrate or simplify the complexity of our subject at the expense of misleading information owners. Our adversaries are always looking for weaknesses in information security, but our strength lies in anticipating sources of threats and having plans in place to prevent the losses that they may cause.

It is impossible to collect a truly complete list of the sources of information losses that can be caused by the intentional or accidental acts of people. We really have no idea what people may do—now or in the future. We base our lists on experience, but until we can conceive of an act, or until a threat actually surfaces or occurs, we cannot include it on the list. And not knowing the threat means that we cannot devise a plan to protect against it. This is one of the reasons that information security is still a folk art rather than a science.

3.4.4 Challenge of Complete Lists. I believe that my lists of physical sources of loss and information losses are complete, but I am always interested in expanding them to include new sources of loss that I may have overlooked.

While I was lecturing in Australia, for example, a delegate suggested that I had omitted an important category. His computer center had experienced an invasion of field mice with a taste for electrical insulation. The intruders proceeded to chew through the computer cables, ruining them. Consequently, I had to add rodents to my list of

sources. I then heard about an incident in San Francisco in which the entire evening shift of computer operations workers ate together in the company cafeteria to celebrate a birthday. Then they all contracted food poisoning, leaving their company without sufficient operations staff for two weeks. I combined the results of these two events into a category named "Living Organisms."

3.5 FUNCTIONS OF INFORMATION SECURITY. The model for information security that I have proposed includes 12 security functions instead of the 3 (prevention, detection, and recovery) included in previous models. These functions describe the activities that information security practitioners and information owners engage in to protect information, as well as the objectives of the security controls that they use. Every control serves one or more of these functions.

Although some security specialists add other functions to the list, such as quality assurance and reliability, I consider these to be outside the scope of information security; other specialized fields deal with them. Reliability is difficult to relate to security except as endangerment when perpetrators destroy the reliability of information and systems, which is a violation of security. Thus, security must preserve a state of reliability but need not necessarily attempt to improve it. Security must protect the auditability of information and systems while, at the same time, security itself must be reliable and auditable. I believe that my security definitions include destruction of the reliability and auditability of information at a high level of abstraction. For example, reliability is reduced when the authenticity of information is put into question by changing it from a correct representation of fact.

Similarly, I do not include such functions as authentication of users and verification in my lists, since I consider these to be control objectives to achieve the 12 functions of information security.

There is a definite logic to the order in which I present the 12 functions in my list. A methodical information security practitioner is likely to apply the functions in this order when resolving security vulnerabilities.

1. Information security must first be independently *audited* in an adversarial manner in order to document its state and to identify its weaknesses and strengths.
2. The practitioner must determine if a security problem can be *avoided* altogether.
3. If the problem cannot be avoided, the practitioner needs to try to *deter* potential abusers or forces from misbehaving.
4. If the threat cannot be avoided or deterred, the practitioner attempts to *detect* its activation.
5. If detection is not assured, then the practitioner tries to *prevent* the act from occurring.
6. If prevention fails and an act occurs, then the practitioner needs to stop it or minimize its harmful effects through *mitigation*.
7. The practitioner needs to determine if *transferring* the responsibility to another individual or department might be more effective at resolving the situation resulting from the attack, or if another party (e.g., an insurer) might be held accountable for the cost of the loss.
8. After a loss occurs, the practitioner needs to *investigate* and search for the individual(s) or force(s) that caused or contributed to the incident as well as for any parties that played a role in it—positively or negatively.
9. When identified, all parties should be *sanctioned or rewarded* as appropriate.

10. After an incident is concluded, the victim needs to *recover* or assist with recovery.
11. The stakeholders should take *corrective* actions to prevent the same type of incident from occurring again.
12. The stakeholders must learn from the experience in order to advance their knowledge of information security and *educate* others.

3.6 SELECTING SAFEGUARDS USING A STANDARD OF DUE DILIGENCE.

Information security practitioners usually refer to the process of selecting safeguards as risk assessment, risk analysis, or risk management. Selecting safeguards based on risk calculations can be a fruitless and expensive process. Although many security experts and associations advocate using risk assessment methods, many organizations ultimately find that using a standard of due diligence (or care) is far superior and more practical. Often one sad experience of using security risk assessment is sufficient to convince information security departments and corporate management of their limitations. Security risk is a function of probability or frequency of occurrence of rare loss events and their impact, and neither is sufficiently measurable or predictable for investment in security. Note that risk applies only to rare events. Events such as computer virus attacks or credit card fraud are occurring continuously and are not risks; they are certainties and can be measured, controlled, and managed.

The standard of due diligence approach is simple and obvious; it is the default process that I recommend and that is commonly used today instead of more elaborate "scientific" approaches. The standard of due diligence approach is recognized and accepted by many legal documents and organizations and is documented in numerous business guides. The 1996 U.S. federal statute on protecting trade secrets (18 USC §1831), for example, states in (3)(a) that the owner of information must take "reasonable measures to keep such information secret" for it to be defined as a trade secret. (See Chapter 45.)

3.7 THREATS, ASSETS, VULNERABILITIES MODEL. Pulling all of the aspects together in one place is a useful way to analyze security threats and vulnerabilities and to create effective scenarios to test real information systems and organizations. The model illustrated in Exhibit 3.1 is designed to help readers do this. Users can outline a scenario or analyze a real case by circling and connecting the appropriate descriptors in each column of the model.

In this version of the model, the Controls column lists only the subject headings of control types; a completed model would contain hundreds of controls. If the model is being used to conduct a review, I suggest that the Vulnerabilities section of the model be renamed to Recommended Controls.

3.8 CONCLUSION. The security framework proposed in this chapter represents an attempt to overcome the dominant technologist view of information security by focusing more broadly on all aspects of security, including the information that we are attempting to protect, the potential sources of loss, the types of loss, the controls that we can apply to avoid loss, the methods for selecting those controls, and our overall objectives in protecting information. This broad focus should have two beneficial effects: advancing information security from a narrow folk art to a broad-based discipline and—most important—helping to reduce many of the losses associated with information, wherever it exists.

EXHIBIT 3.1 Threats, Assets, and Vulnerabilities Model

Threats				Assets	Vulnerabilities (Missing and Deficient Controls)		
Offenders Have/Acquire	**Abuse/ Misuse**	**Methods**	**Losses**	**Assets Lost**	**Control Objectives**	**Controls (Types)**	**Control Guides**
Skills learning technology people **Knowledge** direct indirect **Resources** computer services transport financial **Authority** employment contract ownership possession custodian right other	Errors Omissions Negligence Recklessness Delinquency Civil Disputes Conspiracy Nature Disruption Destruction Theft Privacy Trespass Burglary Larceny Forgery Counterfeiting Smuggling Fraud Scam	**External** heat, cold gases, air water chemical bacteria viruses people animals insects collision collapse shear shake vibrate liquefy flows waves separate slides electric magnets aging radiate Sound Light	**Availability and Utility** destroy damage contaminate deny prolong accelerate delay move misplace convert obscure **Integrity and Authenticity** insert use produce modify replace remove	**Information** spoken printed magnetic electronic optical radio biological **Computer** **Commlines** **Networks** **Facilities** **Buildings** **Transport** **People**	Avoidance Deterrence Prevention Detection Mitigation Sanction Transfer Investigate Recovery Correction	Organization Physical Development Automation Operation Voice Network Access Training Motivation Management Applications Printing Audit Disaster Recovery	Cost effective Due care Complete Consistent Performance Sustain Automatic Tolerated Consequences Override Failsafe Default Instrument Auditable Nonrepudiate Secrecy Universal Independent Unpredictable Tamperproof Compartment

(*continued*)

Exhibit 3.1 Threats, Assets, and Vulnerabilities Model (*Continued*)

Threats				Assets	Vulnerabilities (Missing and Deficient Controls)		
Offenders Have/Acquire	**Abuse/ Misuse**	**Methods**	**Losses**	**Assets Lost**	**Control Objectives**	**Controls (Types)**	**Control Guides**
Motives no intent negligence errors and omissions **Intentional** problem solving gain higher ethic **Extreme Advocacy** social political religious	Embezzlement Bribery Extortion Racketeering Infringement Plagiarism Piracy Espionage Antitrust Contract Securities Employment Kickbacks Laundering Libel Drugs Pornography Harassment Assault Sex attack Kidnapping Murder Suicide	Radio Atomic **Masquerade** impersonate spoof **Programmed** Trojan virus bomb bypass trapdoor **Authority** violation **Active** deny service false data entry **Passive** browse observe **Failure** omit duty **Indirect** crime use	append reorder misrepresent Repudiate Fail to use **Confidential and Possession** locate disclose observe monitor acquire copy take control own infer **Expose to loss** **Endanger** fail instruction				Depth Isolate Least Accountability Trust Multifunction Deception Positional Transparent

Source: Donn B. Parker, *Fighting Computer Crime* (New York, NY: John Wiley & Sons, 1998).

3.9 FURTHER READING

Parker, D. B. *Fighting Computer Crime: A New Framework for Protecting Information*. Wiley (ISBN 978-0471163787), 1998. 528 pp.

Parker, D. B. "What's wrong with information security and how to fix it. Lecture at the Naval Postgraduate School (2005-04-28)." YouTube. www.youtube.com/watch?v=RW9hOBCSy0g

CHAPTER 4

HARDWARE ELEMENTS OF SECURITY

Sy Bosworth and Stephen Cobb

4.1 INTRODUCTION 4·2

4.2 BINARY DESIGN 4·2
- 4.2.1 Pulse Characteristics 4·2
- 4.2.2 Circuitry 4·2
- 4.2.3 Coding 4·3

4.3 PARITY 4·4
- 4.3.1 Vertical Redundancy Checks 4·4
- 4.3.2 Longitudinal Redundancy Checks 4·4
- 4.3.3 Cyclical Redundancy Checks 4·5
- 4.3.4 Self-Checking Codes 4·5

4.4 HARDWARE OPERATIONS 4·6

4.5 INTERRUPTS 4·7
- 4.5.1 Types of Interrupts 4·7
- 4.5.2 Trapping 4·8

4.6 MEMORY AND DATA STORAGE 4·8
- 4.6.1 Main Memory 4·8
- 4.6.2 Read-Only Memory 4·8
- 4.6.3 Secondary Storage 4·9

4.7 TIME 4·10
- 4.7.1 Synchronous 4·10
- 4.7.2 Asynchronous 4·11

4.8 NATURAL DANGERS 4·11
- 4.8.1 Power Failure 4·11
- 4.8.2 Heat 4·11
- 4.8.3 Humidity 4·11
- 4.8.4 Water 4·12
- 4.8.5 Dirt and Dust 4·12
- 4.8.6 Radiation 4·12
- 4.8.7 Downtime 4·12

4.9 DATA COMMUNICATIONS 4·13
- 4.9.1 Terminals 4·13
- 4.9.2 Wired Facilities 4·14
- 4.9.3 Wireless Communications 4·16

4.10 CRYPTOGRAPHY 4·16

4.11 BACKUP 4·17
- 4.11.1 Personnel 4·18
- 4.11.2 Hardware 4·18
- 4.11.3 Power 4·19
- 4.11.4 Testing 4·20

4.12 RECOVERY PROCEDURES 4·20

4.13 MICROCOMPUTER CONSIDERATIONS 4·20
- 4.13.1 Accessibility 4·20
- 4.13.2 Knowledge 4·21
- 4.13.3 Motivation 4·21
- 4.13.4 Opportunity 4·21
- 4.13.5 Threats to Microcomputers 4·21
- 4.13.6 Maintenance and Repair 4·24

4.14 CONCLUSION 4·25

4.15 HARDWARE SECURITY CHECKLIST 4·25

4.16 FURTHER READING 4·27

4.1 INTRODUCTION. Computer hardware has always played a major role in computer security. Over the years, that role has increased dramatically, due to both the increases in processing power, storage capacity, and communications capabilities as well as the decreases in cost and size of components. The ubiquity of cheap, powerful, highly connected computing devices poses significant challenges to computer security. At the same time, the challenges posed by large, centralized computing systems have not diminished. An understanding of the hardware elements of computing is thus essential to a well-rounded understanding of computer security.

Chapter 1 of this *Handbook* has additional history of the evolution of information technology.

4.2 BINARY DESIGN. Although there are wide variations among computer architectures and hardware designs, all have at least one thing in common: They utilize a uniquely coded series of electrical impulses to represent any character within their range. Like the Morse code with its dots and dashes, computer pulse codes may be linked together to convey alphabetic or numeric information. Unlike the Morse code, however, computer pulse trains may also be combined in mathematical operations or data manipulation.

In 1946, Dr. John von Neumann, at the Institute for Advanced Study of Princeton University, first described in a formal report how the binary system of numbers could be used in computer implementations. The binary system requires combinations of only two numbers, 0 and 1, to represent any digit, letter, or symbol and, by extension, any group of digits, letters, or symbols. In contrast, the conventional decimal system requires combinations of 10 different numbers, from 0 to 9, letters from a to z, and a large number of symbols, to convey the same information. Von Neumann recognized that electrical and electronic elements could be considered as having only two states, on and off, and that these two states could be made to correspond to the 0 and 1 of the binary system. If the turning on and off of a computer element occurred at a rapid rate, the voltage or current outputs that resulted would best be described as pulses. Despite 60 years of intensive innovation in computer hardware, and the introduction of some optically based methods of data representation, the nature of these electrical pulses and the method of handling them remain the ultimate measure of a computer's accuracy and reliability.

4.2.1 Pulse Characteristics. Ideally, the waveform of a single pulse should be straight-sided, flat-topped, and of an exactly specified duration, amplitude, and phase relationship to other pulses in a series. It is the special virtue of digital computers that they can be designed to function at their original accuracy despite appreciable degradation of the pulse characteristics. However, errors will occur when certain limits are exceeded, and thus data integrity will be compromised. Because these errors are difficult to detect, it is important that a schedule of preventive maintenance be established and rigidly adhered to. Only in this way can operators detect degraded performance before it is severe enough to affect reliability.

4.2.2 Circuitry. To generate pulses of desirable characteristics, and to manipulate them correctly, requires components of uniform quality and dependability. To lower manufacturing costs, to make servicing and replacement easier, and generally to improve reliability, computer designers try to use as few different types of components as possible and to incorporate large numbers of each type into any one machine.

First-generation computers used as many as 30,000 vacuum tubes, mainly in a half-dozen types of logic elements. The basic circuits were flip-flops, or gates, that produced an output pulse whenever a given set of input pulses was present. However, vacuum tubes generated intense heat, even when in a standby condition. As a consequence, the useful operating time between failures was relatively short.

With the development of solid state diodes and transistors, computers became much smaller and very much cooler than their vacuum-tube predecessors. With advances in logic design, a single type of gate, such as the not-and (NAND) circuit, could replace all other logic elements. The resulting improvements in cost and reliability have been accelerated by the use of monolithic integrated circuits. Not least in importance is their vastly increased speed of operation. Since the meantime between failures of electronic computer circuitry is generally independent of the number of operations performed, it follows that throughput increases directly with speed; speed is defined as the rate at which a computer accesses, moves, and manipulates data. The ultimate limitation on computer speed is the time required for a signal to move from one physical element to another. At a velocity of 299,792,458 meters per second (186,282 miles per second) in vacuum, an electrical signal travels 3.0 meters or 9.84 feet in 10 nanoseconds (0.000.000.01 seconds). If components were as large as they were originally, and consequently as far apart, today's nanosecond computer speeds would clearly be impossible, as would be the increased throughput and reliability now commonplace.

4.2.3 Coding. In a typical application, data may be translated and retranslated automatically into a half-dozen different codes thousands of times each second. Many of these codes represent earlier technology retained for backward compatibility and economic reasons only. In any given code, each character appears as a specific group of pulses. Within each group, each pulse position is known as a *bit*, since it represents either of the two *bi*nary digi*ts*, 0 or 1. Exhibit 4.1 illustrates some of the translations that may be continuously performed as data move about within a single computer.

A *byte* is the name originally applied to the smallest group of bits that could be read and written (accessed or addressed) as a unit. Today a byte is always considered by convention to have 8 bits. In modern systems, a byte is viewed as the storage unit for a single American Standard Code for Information Interchange (ASCII) character, although newer systems such as Unicode, which handle international accented characters

EXHIBIT 4.1 Common Codes for Numeral 5

Code	Bits	Typical Use	Bit Pattern for "5"
Hexadecimal	4	Console switches	0101
Baudot	5	Paper tape	00001
Binary-Coded-Decimal (BCD)	6	Console indicators	000101
Transcode	6	Data transmission	110101
USASCII	7	Data transmission	0110101
EBCDIC	8	Buffer	11110101
EBCDIC, zoned decimal	8	Main memory	11000101
EBCDIC, packed decimal	8	Arithmetic logic unit	01011100
USASCII-8	8	Data transmission	01010101
Hollerith	12	Card reader/punch	000000010000
Binary, halfword	16	Arithmetic logic unit	0000000000000101

and many other symbols, use up to 4 bytes per character. By convention, most people use metric prefixes (kilo-, mega-, giga-, tera-) to indicate collections of bytes; thus *KB* refers to *kilobytes* and is usually defined as 1,024 bytes. Outside the data processing field, *K* would normally indicate the multiplier 1,000. Because of the ambiguity in definitions, the United States National Institute of Standards and Technology (NIST) proposed, and in 2000 the International Electrotechnical Commission established, a new set of units for information or computer storage. These units are established by a series of prefixes to indicate powers of 2; in this scheme, KB means *kibibytes* and refers exclusively to 1,024 (2^{10} or $\sim 10^3$) bytes. However, *kibibytes*, *mebibytes* (2^{20} or $\sim 10^6$), *gibibytes* (2^{30} or $\sim 10^9$), and *tebibytes* (2^{40} or $\sim 10^{12}$) are terms that have not yet become widely used.

Because translations between coding systems are accomplished at little apparent cost, any real incentive to unify the different systems is lacking. However, all data movements and translations increase the likelihood of internal error, and for this reason parity checks and validity tests have become indispensable.

4.3 PARITY. Redundancy is central to error-free data processing. By including extra bits in predetermined locations, certain types of errors can be detected immediately by inspection of these *metadata* (data about the original data). In a typical application, data move back and forth many times, among primary memory, secondary storage, input and output devices, as well as through communications links. During these moves, the data may lose integrity by dropping 1 or more bits, by having extraneous bits introduced, and by random changes in specific bits. To detect some of these occurrences, parity bits are added before data are moved and are checked afterward.

4.3.1 Vertical Redundancy Checks. In this relatively simple and inexpensive scheme, a determination is initially made as to whether there should be an odd or an even number of "1" bits in each character. For example, using the binary-coded decimal representation of the numerical "5," we find that the 6-bit pulse group 000101 contains two 1s, an even number. Adding a seventh position to the code group, we may have either type of parity. If odd parity has been selected, a 1 would be added in the leftmost checkbit position:

Odd parity	1000101 three 1s
Even parity	0000101 two 1s

After any movement the number of 1 bits would be counted, and if not an odd number, an error would be assumed and processing halted. Of course, if 2 bits, or any even number, had been improperly transmitted, no error would be indicated since the number of "1" bits would still be odd.

To compound the problem of nonuniformity illustrated in Exhibit 4.1, each of the 4-, 5-, 6-, 7-, 8-, and 16-bit code groups may have an additional bit added for parity checking. Furthermore, there may be inconsistency of odd or even parity between manufacturers, or even between different equipment from a single supplier.

4.3.2 Longitudinal Redundancy Checks. Errors may not be detected by a vertical redundancy check (VRC) alone, for reasons just discussed. An additional safeguard, of particular use in data transmission and media recording such as tapes and disks, is the longitudinal redundancy check (LRC). With this technique, an extra

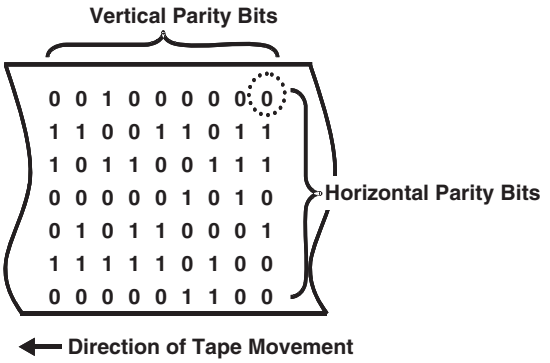


EXHIBIT 4.2 Vertical and Longitudinal Parity, Seven-Track Magnetic Tape

character is generated after some predetermined number of data characters. Each bit in the extra character provides parity for its corresponding row, just as the vertical parity bits do for their corresponding columns. Exhibit 4.2 represents both types as they would be recorded on 7-track magnetic tape. One bit has been circled to show that it is ambiguous. This bit appears at the intersection of the parity row and the parity column, and must be predetermined to be correct for one or the other, as it may not be correct for both. In the illustration, the ambiguous bit is correct for the odd parity requirement of the VRC character column; it is incorrect for the even parity of the LRC bit row.

In actual practice, the vertical checkbits would be attached to each character column as shown, but the longitudinal bits would follow a block of data that might contain 80 to several hundred characters. Where it is possible to use both LRC and VRC, any single data error in a block will be located at the intersection of incorrect row and column parity bits. The indicated bit may then be corrected automatically. The limitations of this method are: (1) multiple errors cannot be corrected, (2) an error in the ambiguous position cannot be corrected, and (3) an error that does not produce both a VRC and LRC indication cannot be corrected.

4.3.3 Cyclical Redundancy Checks. Where the cost of a data error could be high, the added expense of cyclical redundancy checks (CRCs) is warranted. In this technique, a relatively large number of redundant bits is used. For example, each 4-bit character requires 3 parity bits, while a 32-bit computer word needs 6 parity bits. Extra storage space is required in main and secondary memory, and transmissions take longer than without such checks. The advantage, however, is that any single error can be detected, whether in a data bit or a parity bit, and its location can be positively identified. By a simple electronic process of complementation, an incorrect 0 is converted to a 1, and vice versa.

4.3.4 Self-Checking Codes. Several types of codes are in use that inherently contain a checking ability similar to that of the parity system. Typical of these is the 2-of-5 code, in which every decimal digit is represented by a bit configuration containing exactly two 1s and three 0s. Where a parity test would consist of counting 1s to see if their number was odd or even, a 2-of-5 test would indicate an error whenever the number of 1s was more or less than 2.

4.4 HARDWARE OPERATIONS. Input, output, and processing are the three essential functions of any computer. To protect data integrity during these operations, several hardware features are available.

- **Read-after-write**. In disk and tape drives, it is common practice to read the data immediately after they are recorded and to compare them with the original values. Any disagreement signals an error that requires rewriting.
- **Echo**. Data transmitted to a peripheral device, to a remote terminal (see Section 4.9.1), or to another computer can be made to generate a return signal. This echo is compared with the original signal to verify correct reception. However, there is always the risk that an error will occur in the return signal and falsely indicate an error in the original transmission.
- **Overflow**. The maximum range of numerical values that any computer can accommodate is fixed by its design. If a program is improperly scaled, or if an impossible operation such as dividing by zero is called for, the result of an arithmetic operation may exceed the allowable range, producing an overflow error. Earlier computers required programmed instructions to detect overflows, but this function now generally is performed by hardware elements at the machine level. Overflows within application programs still must be dealt with in software. (Indeed, failure to do so can render software susceptible to abuse by malicious parties.)
- **Validation**. In any one computer coding system, some bit patterns may be unassigned, and others may be illegal. In the IBM System/360 Extended Binary Coded Decimal Interchange Code (EBCDIC), for example, the number 9 is represented by 11111001, but 11111010 is unassigned. A parity check would not detect the second group as being in error, since both have the same number of 1 bits. A validity check, however, would reject the second bit configuration as invalid.

  Similarly, certain bit patterns represent assigned instruction codes while others do not. In one computer, the instruction to convert packed-decimal numbers to zoned-decimal numbers is 11110011, or F3 in hexadecimal notation; 11110101, or F5, is unassigned, and a validity check would cause a processing halt whenever that instruction was tested.
- **Replication**. In highly sensitive applications, it is good practice to provide backup equipment on site, for immediate changeover in the event of failure of the primary computer. For this reason, it is sometimes prudent to retain two identical, smaller computers rather than to replace them with a single unit of equivalent or even greater power. Fault-tolerant, or fail-safe, computers use two or more processors that operate simultaneously, sharing the load and exchanging information about the current status of duplicate processes running in parallel. If one of the processors fails, another continues the processing without pause.

  Many sensitive applications, such as airline reservation systems, have extensive data communications facilities. It is important that all of this equipment be duplicated as well as the computers themselves. (The failure of an airline reservation system, if permitted to extend beyond a relatively small number of hours, could lead to failure of the airline itself.)

  Replacements should also be immediately available for peripheral devices. In some operating systems, it is necessary to inform the system that a device is down and to reassign its functions to another unit. In the more sophisticated systems, a malfunctioning device is automatically cut out and replaced. For example, the New York Stock Exchange operates and maintains two identical trading systems so that failure of the primary system should not result in any interruption to trading.

4.5 INTERRUPTS. The sequence of operations performed by a computer system is determined by a group of instructions: a program. However, many events that occur during operations require deviations from the programmed sequence. *Interrupts* are signals generated by hardware elements that detect changed conditions and initiate appropriate action. The first step is immediately to store the status of various elements in preassigned memory locations. The particular stored bit patterns, commonly called program status words, contain the information necessary for the computer to identify the cause of the interrupt, to take action to process it, and then to return to the proper instruction in the program sequence after the interrupt is cleared.

4.5.1 Types of Interrupts. Five types of interrupts are in general use. Each of them is of importance in establishing and maintaining data processing integrity.

4.5.1.1 Input/Output Interrupts. Input/output (I/O) interrupts are generated whenever a device or channel that had been busy becomes available. This capability is necessary to achieve error-free use of the increased throughput provided by buffering, overlapped processing, and multiprogramming.

After each I/O interrupt, a check is made to determine whether the data have been read or written without error. If so, the next I/O operation can be started; if not, an error-recovery procedure is initiated. The number of times that errors occur should be recorded so that degraded performance can be detected and corrected.

4.5.1.2 Supervisor Calls. The supervisor, or monitor, is a part of the operating system software that controls the interactions between all hardware and software elements.

Every request to read or write data is scheduled by the supervisor when called upon to do so. I/O interrupts also are handled by supervisor calls that coordinate them with read/write requests. Loading, executing, and terminating programs are other important functions initiated by supervisor calls.

4.5.1.3 Program Check Interrupts. Improper use of instructions or data may cause an interrupt that terminates the program. For example, attempts to divide by zero and operations resulting in arithmetic overflow are voided. Unassigned instruction codes, attempts to access protected storage, and invalid data addresses are other types of exceptions that cause program check interrupts.

4.5.1.4 Machine Check Interrupts. Among the exception conditions that will cause machine check interrupts are parity errors, bad disk sectors, disconnection of peripherals in use, and defective circuit modules. It is important that proper procedures be followed to clear machine checks without loss of data or processing error.

4.5.1.5 External Interrupts. External interrupts are generated by timer action, by pressing an Interrupt key, or by signals from another computer. When two central processing units are interconnected, signals that pass between them initiate external interrupts. In this way, control and synchronization are continuously maintained while programs, data, and peripheral devices may be shared and coordinated.

In mainframes, an electronic clock generally is included in the central processor unit for time-of-day entries in job logs and for elapsed-time measurements. As an interval timer, the clock can be set to generate an interrupt after a given period. This feature should be used as a security measure, preventing sensitive jobs from remaining on the computer long enough to permit unauthorized manipulation of data or instructions.

4.5.2 Trapping. Trapping is a type of hardware response to an interrupt. Upon detecting the exception, an unconditional branch is taken to some predetermined location. An instruction there transfers control to a supervisor routine that initiates appropriate action.

4.6 MEMORY AND DATA STORAGE. Just as the human mind is subject to aberrations, so is computer memory. In the interests of data security and integrity, various therapeutic measures have been developed for the several types of storage.

4.6.1 Main Memory. Random access memory (RAM), and its derivatives, such as dynamic RAM (DRAM), synchronous DRAM (SDRAM, introduced in 1996 and running at 133 megaHertz [MHz]), and DDR-3 (Double Data Rate 3 SDRAM, announced in 2005 and running at 800 MHz), share the necessary quality of being easily and quickly accessed for reading and writing of data. Unfortunately, this necessary characteristic is at the same time a potential source of difficulty in maintaining data integrity against unwanted read/write operations. The problems are greatly intensified in a multiprogramming environment, especially with dynamic memory allocation, where the possibility exists that one program will write improperly over another's data in memory. Protection against this type of error must be provided by the operating system. Chapter 24 in this *Handbook* discusses operating system security in more detail.

One form of protection requires that main memory be divided into blocks or *pages*; for example, 2,048 eight-bit bytes each. Pages can be designated as read-only when they contain constants, tables, or programs to be shared by several users. Additionally, pages that are to be inaccessible except to designated users may be assigned a lock by appropriate program instructions. If a matching key is not included in the user's program, access to that page will be denied. Protection may be afforded against writing only or against reading and writing.

4.6.2 Read-Only Memory. One distinguishing feature of main memory is the extremely high speed at which data can be entered or read out. The set of sequential procedures that accomplishes this and other functions is the program, and the programmer has complete freedom to combine any valid instructions in a meaningful way. However, certain operations, such as system start-up, or *booting,* are frequently and routinely required, and they may be performed automatically by a preprogrammed group of memory elements. These elements should be protected from inadvertent or unauthorized changes.

For this purpose, a class of memory elements has been developed that, once programmed, cannot be changed at all, or require a relatively long time to do so. These elements are called *read-only memory*, or ROM; the process by which sequential instructions are set into these elements is known as *microprogramming*. The technique can be used to advantage where data integrity is safeguarded by eliminating the possibility of programmer error.

Variations of the principle include programmable read-only memories (PROM) and erasable, programmable read-only memory (EPROM), all of which combine microprogramming with a somewhat greater degree of flexibility than read-only memory itself. The data on these chips can be changed through a special operation often referred to as *flashing* (literally exposure to strong ultraviolet light; this is different from *flash memory* used today for storage of such data as digital music files and digital photographs—we will return to the subject of *flash memory* in the next section).

4.6.3 Secondary Storage. The term "secondary storage" traditionally has been used to describe storage such as magnetic disks, diskettes, tapes, and tape cartridges. Although the 1.44 megabyte (MB) magnetic floppy disk is obsolete, the magnetic hard drive, with capacities up to terabytes, remains an essential element of virtually all computers, and terabyte-capacity external hard drives the size of a paperback book are now available off-the-shelf for a few hundred dollars.

A more recent development are optical drives such as the removable, *compact disc read-only memory* (CD-ROM), originally made available in the early 1980s, which are useful for long-term archival storage of around 700 MB per disc. Hybrid forms of this type exist as well, such as CD-Rs, which can be written to once, and CD-RWs, which accommodate multiple reads and writes. The digital video disc (DVD), or as it has been renamed, the digital versatile disc, was introduced in 1997 and provides capacities ranging from 4.7 gigabytes (GB) per disc up to 30 GB for data archiving. The higher-capacity optical discs use Blu-ray technology introduced in 2002 and can store 25 GB per side; they typically are used for distributing movies, but BD-R (single use) and BD-RE (rewritable) discs hold much potential for generalized data storage.

The newest addition to secondary storage is RAM that simulates hard disks, known as *flash memory*. Derived from electrical EPROMs (EEPROMs) and introduced by Toshiba in the 1980s, this kind of memory now exists in a huge variety of formats, including relatively inexpensive Universal Serial Bus (USB) tokens with storage capacities now in the gigabyte range. These devices appear as external disk drives when plugged into a plug-and-play personal computer. Another flash memory format is small cards, many the size of postage stamps, that can be inserted into mobile phones, cameras, printers, and other devices as well as computers.

Hardware safeguards described earlier, such as redundancy, validity, parity, and read-after-write, are of value in preserving the integrity of secondary storage. These safeguards are built into the equipment and are always operational unless disabled or malfunctioning. Other precautionary measures are optional, such as standard internal labeling procedures for drives, tapes, and disks. Standard internal labels can include identification numbers, record counts, and dates of creation and expiration. Although helpful, external plastic or paper labels on recordable media are not an adequate substitute for computer-generated labels, magnetically inscribed on the medium itself and automatically checked by programmed instructions.

Another security measure sometimes subverted is write-protection on removable media. Hardware interlocks prevent writing to them. These locks should be activated immediately when the media are removed from the system. Failure to do so will cause the data to be destroyed if the same media are improperly used on another occasion.

Hard drives, optical discs, and flash memory cards are classified as direct access storage devices (DASDs). Unlike magnetic tapes with their exclusively sequential processing, DASDs may process data randomly as well as in sequence. This capability is essential to online operations, where it is not possible to sort transactions prior to processing. The disadvantage of direct access is that there may be less control over entries and more opportunity to degrade the system than exists with sequential batch processing.

One possible source of DASD errors arises from the high rotational velocity of the recording medium and, except on head-per-track devices, the movement of heads as well. To minimize this possibility, areas on the recording surface have their addresses magnetically inscribed. When the computer directs that data be read from or into a particular location, the address in main memory is compared with that read from the DASD. Only if there is agreement will the operation be performed.

Through proper programming, the integrity of data can be further assured. In addition to the address check, comparisons can be made on identification numbers or on key fields within each record. Although the additional processing time is generally negligible, there can be a substantial improvement in properly posting transactions.

Several other security measures often are incorporated into DASDs. One is similar to the protection feature in main memory and relies on determining "extents" for each data set. If these extents, which are simply the upper and lower limits of a data file's addresses, are exceeded, the job will terminate.

Another safety measure is necessitated by the fact that defective areas on a disk surface may cause errors undetectable in normal operations. To minimize this possibility, disks should be tested and certified prior to use and periodically thereafter. Further information is provided by operating systems that record the number of disk errors encountered. Reformatting or replacement must be ordered when errors exceed a predetermined level. Many personal computer hard drives now have some form of Self-Monitoring, Analysis, and Reporting Technology (SMART). Evolved from earlier technology such as IBM's Predictive Failure Analysis (PFA) and Intellisafe by computer manufacturer Compaq, and disk drive manufacturers Seagate, Quantum, and Conner, SMART can alert operators to potential drive problems. Unfortunately, the implementation of SMART is not standardized, and its potential for preventive maintenance and failure prediction is often overlooked.

Note that SMART is different from the range of technologies used to protect hard drives from head crashes. A head crash occurs when the component that reads data from the disk actually touches the surface of the disk, potentially damaging it and the data stored on it. Many hard drives have systems in place to withdraw heads from the disk before such contact occurs. These protective measures have reached the point where an active hard drive can be carried around in relative safety as part of a music and video player (e.g., Apple iPod or Microsoft Zune).

4.7 TIME. Within the computer room and in many offices, a wall clock is usually a dominant feature. There is no doubt that this real-time indicator is of importance in scheduling and regulating the functions of people and machines, but the computer's internal timings are more important for security.

4.7.1 Synchronous. Many computer operations are independent of the time of day but must maintain accurate relationships with some common time and with each other. Examples of this synchronism include the operation of gates, flip-flops, and registers, and the transmission of data at high speeds. Synchronism is obtained in various ways. For gates and other circuit elements, electronic clocks provide accurately spaced pulses at a high-frequency rate, while disk and tape drives are maintained at rated speed by servomotor controls based on power-line frequency.

Of all computer errors, the ones most difficult to detect and correct are probably those caused by timing inconsistencies. Internal clocks may produce 1 billion pulses per second (known as 1 gigahertz [GHz]), or more, when the computer is turned on. The loss of even a single pulse, or its random deformation or delay, can cause undetected errors. More troublesome is the fact that even if errors are detected, their cause may not be identified unless the random timing faults become frequent or consistent.

An example of the insidious nature of timing faults is the consequence of electrical power fluctuations when voltage drops below standard. During these power transients, disk drives may slow down; if sectors are being recorded, their physical size will be

correspondingly smaller. Then, when the proper voltage returns, the incorrect sector sizes can cause data errors or loss.

4.7.2 Asynchronous. Some operations do not occur at fixed time intervals and therefore are termed "asynchronous." In this mode, signals generated by the end of one action initiate the following one. As an example, low-speed data transmissions such as those using ordinary modems are usually asynchronous. Coded signals produced by the random depression of keyboard keys are independent of any clock pulses.

4.8 NATURAL DANGERS. To preserve the accuracy and timeliness of computer output, computers must be protected against environmental dangers. Chapters 22 and 23 of this *Handbook* discuss such threats in extensive detail.

4.8.1 Power Failure. Probably the most frequent cause of computer downtime is power failure. Brownouts and blackouts are visible signs of trouble; undetected voltage spikes are far more common, although hardly less damaging.

Lightning can produce voltage spikes on communications and power lines of sufficient amplitude to destroy equipment or, at the very least, to alter data randomly. Sudden storms and intense heat or cold place excessive loads on generators. The drop in line voltage that results can cause computer or peripheral malfunction. Even if it does not, harmful voltage spikes may be created whenever additional generators are switched in to carry higher loads.

Where warranted, a recording indicator may be used to detect power-line fluctuations. Such monitoring often is recommended when computer systems show unexplained, erratic errors. At any time that out-of-tolerance conditions are signaled, the computer output should be checked carefully to ensure that data integrity has not been compromised. If such occurrences are frequent, or if the application is a sensitive one, auxiliary power management equipment should be considered. These range from simple voltage regulators and line conditioners to uninterruptible power supplies (UPSs).

4.8.2 Heat. Sustained high temperatures can cause electronic components to malfunction or to fail completely. Air conditioning (AC) is therefore essential, and all units must be adequate, reliable, and properly installed. If backup electrical power is provided for the computer, it must also be available for the air conditioners. For example, after the San Francisco earthquake of 1989, the desktop computers and network servers in at least one corporate headquarters were damaged by a lack of synchronization between air conditioning and power supply. The AC was knocked out by the quake, and the building was evacuated, but the computers were left on. Many of them failed at the chip and motherboard level over the next few days because the temperature in the uncooled offices got too high. A frequently unrecognized cause of overheating is obstruction of ventilating grilles. Printouts, tapes, books, and other objects must not be placed on cabinets where they can prevent free air circulation. A digital thermometer is a good investment for any room in which computers are used. Today, many electronic devices include thermostats that cut off the power if internal temperatures exceed a danger limit.

4.8.3 Humidity. Either extreme of humidity can be damaging. Low humidity—below about 20 percent—permits buildup of static electricity charges that

may affect data pulses. Because this phenomenon is intensified by carpeting, computer room floors should either be free of carpeting or covered with antistatic carpet.

High humidity—above about 80 percent—may lead to condensation that causes shorts in electrical circuits or corrodes metal contacts. To ensure operation within acceptable limits, humidity controls should be installed and a continuous record kept of measured values.

4.8.4 Water. Water introduced by rain, floods, bursting pipes, and overhead sprinklers probably has been responsible for more actual computer damage than fire or any other single factor. Care taken in locating computer facilities, in routing water pipes, and in the selection of fire-extinguishing agents will minimize this significant danger.

The unavailability of water—following a main break, for example—will cause almost immediate shutdown of water-cooled mainframes. Mission-critical data centers should be prepared for this contingency. As an example, when the Des Moines River flooded in 1993, it caused the skyscraper housing the headquarters of the Principal Financial Group to be evacuated, but not because of water in the building. The building stayed high and dry, but the flood forced the city water plant to shut down, depriving the building of the water necessary for cooling. After the flood, the company installed a 40,000-gallon water tank in the basement, to prevent any recurrence of this problem.

4.8.5 Dirt and Dust. Particles of foreign matter can interfere with proper operation of magnetic tape and disk drives, printers, and other electromechanical devices. All air intakes must be filtered, and all filters must be kept clean. Cups of coffee seem to become especially unstable in a computer environment; together with any other food or drink, they should be banned entirely. Throughout all areas where computer equipment is in use, good housekeeping principles should be rigorously enforced.

4.8.6 Radiation. Much has been written about the destructive effect of magnetic fields on tape or disk files. However, because magnetic field strength diminishes rapidly with distance, it is unlikely that damage actually could be caused except by large magnets held very close to the recorded surfaces. For example, storing a CD or DVD by attaching it to a filing cabinet with a magnet is not a good idea, but simply walking past a refrigerator decorated with magnets while holding a CD or DVD is unlikely to do any damage.

The proliferation of wireless signals can expose data to erroneous pulses. Offices should be alert for possible interference from and between cordless phones, mobile phones, wireless Internet access points and peripherals, and microwave ovens.

Radioactivity may be a great threat to personnel but not to the computer or its recording media.

4.8.7 Downtime. It is essential to the proper functioning of a data center that preventive maintenance be performed regularly and that accurate records be kept of the time and the reason that any element of the computer is inoperative. The more often the computer is down, the more rushed operators will be to catch up on their scheduled workloads. Under such conditions, controls are bypassed, shortcuts are taken, and human errors multiply.

Downtime records should be studied to detect unfavorable trends and to pinpoint equipment that must be overhauled or replaced before outages become excessive.

If unscheduled downtime increases, preventive maintenance should be expanded or improved until the trend is reversed.

4.9 DATA COMMUNICATIONS. One of the most dynamic factors in current computer usage is the proliferation of devices and systems for data transmission. These range from telephone modems to wired networks, from Internet-enabled cell phones to 802.11 wireless Ethernet, and include Bluetooth, infrared, personal digital assistants (PDAs), music players, and new technologies that appear almost monthly. Computers that do not function at least part time in a connected mode may well be rarities. For fundamentals of data communications, see Chapter 5 of this *Handbook*.

The necessity for speeding information over great distances increases in proportion to the size and geographic dispersion of economic entities; the necessity for maintaining data integrity and security, and the difficulty of doing so, increases even more rapidly. Major threats to be guarded against include human and machine errors, unauthorized accession, alteration, and sabotage. The term "accession" refers to an ability to read data stored or transmitted within a computer system; it may be accidental or purposeful. "Alteration" is the willful entering of unauthorized or incorrect data. "Sabotage" is the intentional act of destroying or damaging the system or the data within it. For each of these threats, the exposure and the countermeasures will depend on the equipment and the facilities involved.

4.9.1 Terminals. In these discussions, a *terminal* is any input/output device that includes facilities for receiving, displaying, composing, and sending data. Examples include personal computers and specialized devices such as credit card validation units.

Data communications are carried on between computers, between terminals, or between computers and terminals. The terminals themselves may be classified as *dumb* or *intelligent*. Dumb terminals have little or no processing or storage capability and are largely dependent on a host computer for those functions. Intelligent terminals generally include disk storage and capabilities roughly equivalent to those of a personal computer. In addition to vastly improved communications capabilities, they are capable of stand-alone operation.

In the simplest of terminals, the only protection against transmission errors lies in the inability to recognize characters not included in the valid set and to display a question mark or other symbol when one occurs. Almost any terminal can be equipped to detect a vertical parity error. More sophisticated terminals are capable of detecting additional errors through longitudinal and cyclical redundancy characters, as well as by vertical parity and validity checks. Of course, error detection is only the first step in maintaining data integrity. Error correction is by far the more important part, and retransmission is the most widely used correction technique.

Intelligent terminals and personal computers are capable of high-speed transmission and reception. They can perform complicated tests on data before requesting retransmission, or they may even be programmed to correct errors internally. The techniques for self-correction require forward-acting codes, such as the Hamming cyclical code. These are similar to the error-detecting cyclic redundancy codes, except that they require even more redundant bits. Although error correction is more expensive and usually slower than detection with retransmission, it is useful under certain circumstances. Examples include simplex circuits where no return signal is possible, and half-duplex circuits where the time to turn the line around from transmission to reception is too long. Forward correction is also necessary where errors are so

numerous that retransmissions would clog the circuits, with little or no useful information throughput.

A more effective use of intelligent terminals and personal computers is to preserve data integrity by encryption, as described in this chapter and in Chapter 7. Also, they may be used for compression or compaction. Reducing the number of characters in a message reduces the probability of an error as well as the time required for transmission. One technique replaces long strings of spaces or zeroes with a special character and a numerical count; the procedure is reversed when receiving data.

Finally, the intelligent terminal or microprocessor may be used to encode or decipher data when the level of security warrants cryptography.

All terminals, of every type, including desktop and notebook personal computers (PCs), have at least one thing in common: the need to be protected against sabotage or unauthorized use. Although the principles for determining proper physical location and the procedures for restricting access are essentially the same as those that apply to a central computer facility, the actual problems of remote terminals are even more difficult. Isolated locations, inadequate supervision, and easier access by more people all increase the likelihood of compromised security.

4.9.2 Wired Facilities. Four types of wired facilities are in widespread use: dial-up access, leased lines, digital subscriber lines (DSL), and cable carriers. Both common carriers and independent systems may employ various media for data transmission. The increasing need for higher speed and better quality in data transmission has prompted utilization of coaxial and fiber optic cables, while microwave stations and communication satellites often are found as wireless links within wired systems.

Generally, decisions as to the choice of service are based on the volume of data to be handled and on the associated costs, but security considerations may be even more important.

4.9.2.1 Dial-Up Lines. Still widely used for credit and debit card terminals, dial-up lines have been replaced for many other applications by leased lines, DSL lines, and cables carrying Internet traffic (using the TCP/IP protocol discussed in Chapter 5 of this *Handbook*). Dial-up connections are established between modems operating over regular voice lines sometimes referred to as *plain old telephone service* (POTS).

Where dial-up access to hardware still exists, for example, for maintenance of certain equipment, proper controls are essential to protect both the equipment and the integrity of other systems to which it might be connected. Dial-up ports may be reached by anyone with a phone, anywhere on the planet, and the practice of *war-dialing* to detect modems is still used by those seeking unauthorized access to an organization's network. (War dialing involves dialing blocks of numbers to find which ones respond as modems or fax machines. These numbers are recorded and may be dialed later in an attempt to gain unauthorized access to systems or services.) It is advisable to:

- Compile a log of unauthorized attempts at entry, and use it to discourage further efforts.
- Compile a log of all accesses to sensitive data, and verify their appropriateness.
- Equip all terminals with internal identification generators or answer-back units, so that even a proper password would be rejected if sent from an unauthorized

terminal. This technique may require the availability of an authorized backup terminal in the event of malfunction of the primary unit.

- Provide users with personal identification in addition to a password if the level of security requires it. The additional safeguard could be a magnetically striped or computerized plastic card to be inserted into a special reader. The value of such cards is limited, since they can be used by anyone, whether authorized or not. For high-security requirements, other hardware-dependent biometric identifiers, such as handprints and voiceprints, should be considered.
- Where appropriate, utilize call-back equipment that prevents a remote station from entering a computer directly. Instead, the device dials the caller from an internal list of approved phone numbers to make the actual connection.

With proper password discipline, problems of accession, alteration, and data sabotage can be minimized. However, the quality of transmissions is highly variable. Built into the public telephone system is an automatic route-finding mechanism that directs signals through uncontrollable paths. The distance and the number of switching points traversed, and the chance presence of cross-talk, transients, and other noise products will have unpredictable effects on the incidence of errors. Parity systems, described earlier, are an effective means of reducing such errors.

4.9.2.2 Leased Lines. Lines leased from a common carrier for the exclusive use of one subscriber are known as *dedicated lines*. Because they are directly connected between predetermined points, normally they cannot be reached through the dial-up network. Traditionally, leased lines were copper, but point-to-point fiber optic and coaxial cable lines can also be leased.

Wiretapping is a technically feasible method of accessing leased lines, but it is more costly, more difficult, and less convenient than dialing through the switched network. Leased lines are generally more secure than those that can be readily war-dialed.

To this increased level of security for leased lines is added the assurance of higher-quality reception. The problems of uncertain transmission paths and switching transients are eliminated, although other error sources are not. In consequence, parity checking remains a minimum requirement.

4.9.2.3 Digital Subscriber Lines. Falling somewhere in between a leased line and POTS, a digital subscriber line offers digital transmission locally over ordinary phone lines that can be used simultaneously for voice transmission. This is possible because ordinary copper phone lines can carry, at least for short distances, signals that are in a much higher-frequency range than the human voice. A DSL modem is used by a computer to reach the nearest telephone company switch, at which point the data transmission enters the Internet backbone. Computers connected to the Internet over DSL communicate using TCP/IP and are said to be hosts rather than terminals. They are prone to compromise through a wide range of exploits. However, few if any of these threats are enabled by the DSL itself. As with leased lines, wiretapping is possible, but other attacks, such as exploiting weaknesses in TCP/IP implementations on host machines, are easier.

4.9.2.4 Cable Carriers. Wherever cable television (TV) is available, the same optical fiber or coaxial cables that carry the TV signal also can be used to provide high-speed data communications. The advantages of this technology include download

speeds that can, in the case of coaxial cables, exceed 50 megabits per second, or in the case of fiber optic cable, exceed 100 gigabits per second.

The disadvantages arise from the fact that connections to the carrier may be shared by other subscribers in the same locality. Unless the service provider limits access, perhaps in accordance with a quality-of-service agreement, multiple subscribers can be online simultaneously and thus slow down transmission speeds. Even more serious is the possibility of security breaches, since multiple computers within a neighborhood may be sharing part of a virtual local area network, and thus each is potentially accessible to every other node on that network. For this reason alone, cable connections should be firewalled. For details of firewalls and their uses, see Chapter 26 in this *Handbook*. Another reason for using firewalls is that cable connections are always on, providing maximal opportunity for hackers to access an unattended computer.

4.9.3 Wireless Communications. Data transfers among multinational corporations have been growing very rapidly, and transoceanic radio and telephone lines have proved too costly, too slow, too crowded, and too error-prone to provide adequate service. An important alternative is the communications satellite. Orbiting above Earth, the satellite reflects ultra-high-frequency radio signals that can convey a television program or computer data with equal speed and facility.

For communications over shorter distances, the cost of common-carrier wired services has been so high as to encourage competitive technologies. One of these, the microwave radio link, is used in many networks. One characteristic of such transmissions is that they can be received only on a direct line-of-sight path from the transmitting or retransmitting antenna. With such point-to-point ground stations, it is sometimes difficult to position the radio beams where they cannot be intercepted; with satellite and wireless broadcast communications, it is impossible. This is a significant issue with wireless local area network technology based on the IEEE 802.11 standards and commonly known as Wi-Fi (a brand name owned by the Wi-Fi Alliance; the term is short for *wi*reless *fi*delity). The need for security is consequently greater, and scramblers or cryptographic encoders are essential for sensitive data transfers.

Because of the wide bandwidths at microwave frequencies, extremely fast rates of data transfer are possible. With vertical, longitudinal, and cyclical redundancy check characters, almost all errors can be detected, yet throughput remains high.

4.10 CRYPTOGRAPHY. Competitive pressures in business, politics, and international affairs continually create situations where morality, privacy, and the laws all appear to give way before a compelling desire for gain. Information, for its own sake or for the price it brings, is an eagerly sought after commodity. We are accustomed to the sight of armored cars and armed guards transporting currency, yet often invaluable data are moved with few precautions. When the number of computers and competent technicians was small, the risk in careless handling of data resources was perhaps not great. Now, however, a very large population of knowledgeable computer people exists, and within it are individuals willing and able to use their knowledge for illegal ends. Others find stimulation and satisfaction in meeting the intellectual challenge that they perceive in defeating computer security measures.

Acquiring information in an unauthorized manner is relatively easy when data are communicated between locations. One method of discouraging this practice, or rendering it too expensive to be worth the effort, is cryptographic encoding of data prior to transmission. This technique is also useful in preserving the security of files within data storage devices. If all important files were stored on magnetic or optical media

in cryptographic cipher only, the incidence of theft and resale would unquestionably be less.

Many types of ciphers might be used, depending on their cost and the degree of security required. Theoretically, any code can be broken, given enough time and equipment. In practice, if a cipher cannot be broken fairly quickly, the encoded data are likely to become valueless. However, since the key itself can be used to decipher later messages, it is necessary that codes or keys be changed frequently.

For further information on cryptography, refer to Chapter 7 in this *Handbook*.

4.11 BACKUP. As with most problems, the principal focus in computer security ought to be on prevention rather than on cure. No matter how great the effort, however, complete success can never be guaranteed. There are four reasons for this being so:

1. Not every problem can be anticipated.
2. Where the cost of averting a particular loss exceeds that of recovery, preventive measures may not be justified.
3. Precautionary measures, carried to extremes, can place impossible constraints on the efficiency and productivity of an operation. It may be necessary, therefore, to avoid such measures aimed at events whose statistical probability of occurrence is small.
4. Even under optimum conditions, carefully laid plans may go astray. In the real world of uncertainty and human fallibility, where there is active or inadvertent interference, it is almost a certainty that at one time or another, the best of precautionary measures will prove to be ineffective.

Recognizing the impossibility of preventing all undesired actions and events, it becomes necessary to plan appropriate means of recovering from them. Such plans must include backup for personnel, hardware, power, physical facilities, data, and software. Data backups are discussed more fully in Chapter 57 of this *Handbook*.

Responding to emergencies is described in Chapters 56 of this *Handbook* and business continuity planning and disaster recovery are discussed in Chapter 58 and 59.

Backup plans should be evaluated with respect to:

- The priorities established for each application, to ensure that they are properly assigned and actually observed.
- The time required to restore high-priority applications to full-functioning status.
- The degree of assurance that plans actually can be carried out when required. For important applications, alternative plans should be available in the event that the primary plan cannot be implemented.
- The degree of security and data integrity that will exist if backup plans actually are put into effect.
- The extent to which changing internal or external conditions are noted, and the speed with which plans are modified to reflect such changes.

The assignment of priorities in advance of an actual emergency is an essential and critically important process. In most organizations, new applications proliferate, while old ones are rarely discarded. If backup plans attempt to encompass all jobs,

they are likely to accomplish none. Proper utilization of priorities will permit realistic scheduling, with important jobs done on time and at acceptable costs.

4.11.1 Personnel. The problems of everyday computer operation require contingency plans for personnel on whose performance hardware functioning depends. Illnesses, vacations, dismissals, promotions, resignations, overtime, and extra shifts are some of the reasons why prudent managers are continuously concerned with the problem of personnel backup. The same practices that work for everyday problems can provide guidelines for emergency backup plans. This subject is covered more fully in Chapter 45 of this *Handbook*.

4.11.2 Hardware. Hardware backup for data centers can take several forms:

- Multiple processors at the same site to protect against loss of service due to breakdown of one unit
- Duplicate installations at nearby facilities of the same company
- Maintaining programs at a compatible service bureau, on a test or standby basis
- A contract for backup at a facility dedicated to disaster recovery
- A reciprocal agreement with a similar installation at another company

The probability of two onsite processors both being down at the same time due to internal faults is extremely small. Consequently, most multiple installations rarely fall behind on mission-critical applications. However, this type of backup offers no protection against power failure, fire, vandalism, or any disaster that could strike two or more processors at once. The disasters of September 11, 2001, proved that even a highly unlikely event actually could occur. With duplicate processors at different but commonly owned sites, there is little chance of both being affected by the same forces. Although the safety factor increases with the distance separating them, the difficulty of transporting people and data becomes greater. An alternate site must represent a compromise between these conflicting objectives. Furthermore, complete compatibility of hardware and software will have to be preserved, even though doing so places an undue operational burden on one of the installations. Shortly after September 11, a number of New York financial firms were back in operation with their alternative computer sites across the Hudson River.

The backup provided by service bureaus can be extremely effective, particularly if the choice of facility is carefully made. Although progressive service bureaus frequently improve both hardware and software, they almost never do so in a way that would cause compatibility problems for their existing customers. Once programs have been tested, they can be stored offline on tape or disk at little cost. Updated masters can be rotated in the service bureau library, providing offsite data backup as well as the ability to become fully operational at once.

Effective hardware backup is also available at independent facilities created expressly for that purpose. In one type of facility, there are adequate space, power, air conditioning, and communication lines to accommodate a very large system. Most manufacturers are able to provide almost any configuration on short notice when disaster strikes a valued customer. The costs for this type of base standby facility are shared by a number of users so that expenses are minimal until an actual need arises. However, if two or more sharers are geographically close, their facilities may be rendered inoperative by the same fire, flood, or power failure. Before contracting for such a facility,

it is necessary to analyze this potential problem; the alternative is likely to be a totally false sense of security. Several firms whose facilities were damaged or destroyed on September 11 were provided with complete replacement equipment by their vendors within a short time.

Another type of backup facility is already equipped with computers, disk and tape drives, printers, terminals, and communications lines so that it can substitute instantly for an inoperative system. The standby costs for this service are appreciably more than for a base facility, but the assurance of recovery in the shortest possible time is far greater. Here, too, it would be prudent to study the likelihood of more than one customer requiring the facility at the same time and to demand assurance that one's own needs will be met without fail. Several companies successfully availed themselves of this type of backup and disaster recovery after September 11.

Backup by reciprocal agreement was for many years an accepted practice, although not often put to the test. Unfortunately, many managers still rely on this outmoded safeguard. One has only to survive a single major change of operating system software to realize that when it occurs, neither the time nor the inclination is available to modify and test another company's programs. Even the minor changes in hardware and software that continuously take place in most installations could render them incompatible. At the same time, in accordance with Parkinson's Law, workloads always expand to fill the available time and facilities. In consequence, many who believe that they have adequate backup will get little more than an unpleasant surprise, should they try to avail themselves of the privilege.

4.11.3 Power. The one truly indispensable element of any data processing installation is electric power. Backing up power to PCs and small servers by uninterruptible power supplies is reasonable in cost and quite effective. For mainframes and large servers, several types of power backup are available. The principal determinant in selection should be the total cost of anticipated downtime and reruns versus the cost of backup to eliminate them. Downtime and rerun time may be extrapolated from records of past experience.

Problems due to electrical power may be classified by type and by the length of time that they persist. Power problems as they affect computers consist of variations in amplitude, frequency, and waveform, with durations ranging from fractions of a millisecond to minutes or hours. Long-duration outages usually are due to high winds, ice, lightning, vehicles that damage power lines, or equipment malfunctions that render an entire substation inoperative. For mainframes in data centers, it is usually possible, although costly, to contract for power to be delivered from two different substations, with one acting as backup.

Another type of protection is afforded by gasoline or diesel motor generators. Controls are provided that sense a power failure and automatically start the motor. Full speed is attained in less than a minute, and the generator's output can power a computer for days if necessary.

The few seconds' delay in switching power sources is enough to abort programs running on the computer and to destroy data files. To avoid this, the "uninterruptible" power supply was designed. In one version, the AC power line feeds a rectifier that furnishes direct current to an inverter. The inverter in turn drives a synchronous motor coupled to an alternator whose AC output powers the computer. While the rectifier is providing DC to the inverter, it also charges a large bank of heavy-duty batteries. As soon as a fault is detected on the main power line, the batteries are instantaneously and automatically switched over to drive the synchronous motor. Because the huge drain

on the batteries may deplete them in a few minutes, a diesel generator must also be provided. The advantages of this design are:

- Variations in line frequency, amplitude, and waveform do not get through to the computer.
- Switchover from power line to batteries is undetectable by the computer. Programs keep running, and no data are lost.
- Millisecond spikes and other transients that may be responsible for equipment damage, and undetected data loss are completely suppressed.

A fuller treatment of physical threats is presented in Chapters 22 and 23 of this *Handbook*.

4.11.4 Testing. The most important aspect of any backup plan is its effectiveness. Will it work? It would be a mistake to wait for an emergency to find out. The only sensible alternative is systematic testing.

One form of test is similar to a dress rehearsal, with the actual emergency closely simulated. In this way the equipment, the people, and the procedures can all be exercised, until practice assures proficiency. Periodically thereafter the tests should be repeated, so that changes in hardware, software, and personnel will not weaken the backup capability.

4.12 RECOVERY PROCEDURES. The procedures required to recover from any system problem will depend on the nature of the problem and on the backup measures that were in place. Hardware recovery ranges from instantaneous and fully automatic, through manual repair or replacement of components, to construction, equipping, and staffing of an entirely new data center. Chapters 58 and 59 of this *Handbook* provide extensive information about these issues.

Almost every data center is a collection of equipment, with options, modifications, additions, and special features. Should it become necessary to replace the equipment, a current configuration list must be on hand and the procedures for reordering established in advance. An even better practice would be to keep a current list of *desired* equipment that could be used as the basis for replacement. Presumably, the replacements would be faster and more powerful, but additional time should be scheduled for training and conversion.

4.13 MICROCOMPUTER CONSIDERATIONS. Four factors operate to intensify the problems of hardware security as they relate to small computers:

1. Accessibility
2. Knowledge
3. Motivation
4. Opportunity

4.13.1 Accessibility. *Accessibility* is a consequence of operating small computers in a wide-open office environment rather than in a controlled data center. No security guards, special badges, man-traps, cameras, tape librarians, or shift supervisors limit access to equipment or data media in the office, as they do in a typical data center.

4.13.2 Knowledge. *Knowledge* and its lack are equally dangerous. On one hand, as personal computers pervade the office environment, technical knowledge becomes widely disseminated. Where once this knowledge was limited to relatively few computer experts who could be controlled rather easily, its growing universality now makes control extremely difficult, if not impossible. On the other hand, when computers are operated by people with minimal knowledge and skills, the probability of security breaches through error and inadvertence is greatly increased.

4.13.3 Motivation. *Motivation* exists in numerous forms. It is present wherever valuable assets can be diverted for personal gain; it arises when real or fancied injustice creates a desire for revenge; and it can simply be a form of self-expression.

The unauthorized diversion of corporate assets always has provided opportunities for theft; now, with many employees owning computers at home, the value of stolen equipment, programs, and data can be realized without the involvement of third parties. When a third party is added to the equation and the thriving market in purloined personal data is factored in, the potential for data theft, a low-risk/high-return crime, is greatly increased.

Computers and networks are also a target for sabotage as well as data theft. The reliance upon such systems by governments, the military, large corporations, and other perceived purveyors of social or economic ills means that criminal acts are likely to continue. Because personal computers are now part of these systems, they are also a link to any policy or practice of which one or more groups of people disapprove. The motivation for sabotaging personal computers is more likely in the near term to increase than it is to disappear.

A third motivation for breaching computer security is the challenge and excitement of doing so. Whether trying to overcome technical hurdles, to break the law with impunity, or merely to trespass on forbidden ground, some hackers find these challenges irresistible, and they become criminal hackers. To view such acts with amused tolerance or even mild disapproval is totally inconsistent with the magnitude of the potential damage and the sanctity of the trust barriers that are crossed. Since the technology exists to lock out all but the most determined and technically proficient criminal hacker, failure to protect sensitive systems is increasingly viewed as negligence.

4.13.4 Opportunity. With so many personal computers in almost every office, with virtually no supervision during regular hours, and certainly none at other times, opportunities are plentiful for two types of security breaches: intentional by those with technical knowledge and unintentional by those without.

4.13.5 Threats to Microcomputers. Among the most significant threats to microcomputers are those pertaining to:

- Physical damage
- Theft
- Electrical power
- Static electricity
- Data communications
- Maintenance and repair

4.13.5.1 Physical Damage. Microcomputers and their peripheral devices are not impervious to damage. Disk drives are extremely susceptible to failure through impact; keyboards cannot tolerate dirt or rough handling. It is essential that computers be recognized as delicate instruments and that they be treated accordingly.

Even within an access-controlled data center, where food and drinks are officially banned, it is not uncommon for a cup of coffee to be spilled when set on or near operating equipment. In an uncontrolled office environment, it is rare that one does not see personal computers in imminent danger of being doused with potentially damaging liquids. The problem is compounded by the common practice of leaving unprotected media such as CDs and DVDs lying about on the same surface where food and drink could easily reach them. Although it may not be possible to eliminate these practices entirely, greater discipline will protect data media and equipment from contamination.

As mentioned in the section on heat, damage also can result from blocking vents necessary for adequate cooling. Such vents can be rendered ineffective by placing the equipment too close to a wall or, in the case of laptops, on soft surfaces, such as carpets, that block vents on the base of the machine. Vents on top of computer housings and cathode ray tube–style displays are too often covered by papers or books that prevent a free flow of cooling air. As a result, the internal temperature of the equipment increases, so that marginal components malfunction, intermittent contacts open, errors are introduced, and eventually the system malfunctions or halts.

4.13.5.2 Theft. The opportunities for theft of personal computers and their data media are far greater than for their larger counterparts. Files containing proprietary information or expensive programs are easily copied to removable media as small as a postage stamp and taken from the premises without leaving a trace. External disk drives are small enough to be carried out in a purse or an attaché case, and new thumb-size USB drives look like key fobs to the uninitiated. (For more information about removable, miniaturized, file storage, see Chapter 1 in this *Handbook*.) The widespread practice of taking portable computers home for evening or weekend work eventually renders even the most conscientious guards indifferent. In offices without guards, the problem is even more difficult. Short of instituting a police state of perpetual surveillance, what is to be done to discourage theft? Equipment can be chained or bolted to desks, or locked within cabinets built for the purpose. Greater diligence in recording and tracking serial numbers, more frequent inventories, and a continuing program of education can help. Most of all, it is essential that the magnitude of the problem be recognized at a sufficiently high management level so that adequate resources are applied to its solution. Otherwise, there will be a continuing drain of increasing magnitude on corporate profitability.

4.13.5.3 Power. Even in a controlled data center, brownouts, blackouts, voltage spikes, sags and surges, and other electrical power disturbances represent a threat. The situation is much worse in a typical office, where personal computers are plugged into existing outlets with little or no thought to the consequences of bad power.

Some of the rudimentary precautions that should be taken are:

- Eliminating, or at least controlling, the use of extension cords, cube taps, and multiple outlet strips. Each unit on the same power line may reduce the voltage available to all of the others, and each may introduce noise on the line.

- Providing line voltage regulators and line conditioners where necessary to maintain power within required limits.
- Banning the use of vacuum cleaners or other electrical devices plugged into the same power line as computers or peripheral devices. Such devices produce a high level of electrical noise, in addition to voltage sags and surges.
- Connecting all ground wires properly. This is especially important in older offices equipped with two-prong outlets that require adapter plugs. The third wire of the plug must be connected to a solid earth ground for personnel safety, as well as for reduction of electrical noise.

In addition, the use of UPSs is highly recommended for all computers and ancillary equipment. These devices are available in capacities from about 200 watts for PCs to virtually unlimited sizes for mainframes. While the power line is operational, a UPS is capable of conditioning the line by removing electrical noise, sags, spikes, and surges. When line voltage drops below a preset value, or when power is completely lost, the UPS converts DC from its internal batteries to the AC required to supply the associated equipment.

Depending on its rating and the load, the UPS may provide standby power for several minutes to several hours. This is enough time to shut down a computer normally, or in the case of large installations, to have a motor generator placed online.

The services of a qualified electrician should be utilized wherever there is a possibility of electric power problems.

4.13.5.4 Static Electricity. After one walks across a carpeted floor on a dry day, the spark that leaps from fingertip to computer may be mildly shocking to a person, but to the computer it can cause serious loss of memory, degradation of data, and even component destruction. These effects are even more likely when people touch components inside a computer without proper grounding.

To prevent this, several measures are available:

- Use a humidifier to keep the humidity above 20 percent relative.
- Remove ordinary carpeting. Replace, if desired, with static-free types.
- Use an antistatic mat beneath chairs and desks.
- Use a grounding strip near each keyboard.
- Wear a grounding bracelet when installing or repairing the components of any electronic equipment.

Touching the grounding strip before operating the computer will drain any static electricity charge through the attached ground wire, as will spraying the equipment periodically with an antistatic spray.

Some combination of these measures will protect personnel, equipment, and data from the sometimes obscure, but always real, dangers of static electricity.

4.13.5.5 Data Communications. Although personal computers perform significant functions in a stand-alone mode, their utility is greatly enhanced by communications to mainframes, to information utilities, and to other small computers, remotely via phone lines or the Internet, or through local area networks. All of the security

issues that surround mainframe communications apply to personal computers, with added complications.

Until the advent of personal computers, almost all terminals communicating with mainframes were "dumb." That is, they functioned much like teletype machines, with the ability only to key in or print out characters, one at a time. In consequence, it was much more difficult to breach mainframe security, intentionally or accidentally, than it is with today's fully intelligent personal computers.

The image of thousands of dedicated hackers dialing up readily available computer access numbers or probing Internet addresses, for illicit fun and games or for illegal financial gain, is no less disturbing than it is real. Countermeasures are available, including:

- Two-way encryption (see Chapter 7)
- Frequent password changes (see Chapter 28)
- Automatic call-back before logging on
- Investigation of unsuccessful logons
- Monitoring of hackers' bulletin boards (see Chapters 12 and 15)
- Firewalls to restrict traffic into and out of the computer (see Chapter 26)
- Antivirus software (see Chapter 41)

Legislation that makes directors and senior officers personally liable for any corporate losses that could have been prevented should have a marked effect on overcoming the current inertia. Prudence dictates that preventive action be taken before, rather than corrective action after, such losses are incurred.

4.13.6 Maintenance and Repair. A regular program of preventive maintenance should be observed for every element of a personal computer system. This should include scheduled cleaning of disk drives and their magnetic heads, keyboards, and printers. A vital element of any preventive maintenance program is the frequent changing of air filters in every piece of equipment. If this is not done, the flow of clean, cool air will be impeded, and failure will almost surely result.

Maintenance options for personal computers, in decreasing order of timeliness, include:

- Onsite management by regular employees
- Onsite maintenance by third parties under an annual agreement
- On-call repair, with or without an agreement
- Carry-in service
- Mail-in service

As personal computers are increasingly applied to functions that affect the very existence of a business, their maintenance and repair will demand more management attention. Redundant equipment and onsite backup will always be effective, but the extended time for offsite repairs will no longer be acceptable. For most business applications, "loaners" or "swappers" should be immediately available, so that downtime will be held to an absolute minimum. Management must assess the importance of each functioning personal computer and select an appropriate maintenance and repair policy.

Accessibility, knowledge, motivation, and opportunity are the special factors that threaten every personal computer installation. Until each of these factors has been addressed, no system can be considered secure.

4.14 CONCLUSION. This chapter has dealt principally with the means by which hardware elements of a data processing system affect the security and integrity of its operations. Many safeguards are integral parts of the equipment itself; others require conscious effort, determination, and commitment.

An effective security program—one that provides both decreased likelihood of computer catastrophe and mitigation of the consequences of damage—cannot be designed or implemented without considerable expenditures of time and money. As with other types of loss avoidance, the premium should be evaluated against the expected costs. Once a decision has been made, however, this equivalent to an insurance policy should not be permitted to lapse. The premiums must continue to be paid in the form of periodic testing, continuous updating, and constant vigilance.

For more detailed information about risk management, see Chapter 62 in this *Handbook*. For a discussion of insurance policies against information systems disasters of all kinds, see Chapter 60.

4.15 HARDWARE SECURITY CHECKLIST

Mainframes

- Are security and integrity requirements considered when selecting new equipment?
- Is a schedule of preventive maintenance enforced?
- Is a log kept of all computer malfunctions and unscheduled downtime?
- Is there an individual with responsibility for reviewing the log and initiating action?
- Are parity checks used wherever possible?
- Is there an established procedure for recording parity errors and recovering from them?
- Are forward-acting or error-correcting codes used when economically justified?
- Do operators follow prescribed procedures after a read error or other machine check halt?
- Are all operator interventions logged and explained?
- Is a job log maintained, and is it compared regularly with an authorized run list?
- Is the interval timer used to prevent excessively long runs?
- Are storage protect features such as data locks and read-only paging used?
- Are keys to software data locks adequately protected?
- Are precautions taken to prevent loss of data from volatile memory during power interruptions?
- Are standard internal and external tape and disk labeling procedures enforced?
- Are write-enable protection rings always removed from tape reels immediately after use?
- Is there a rule that new tapes and disks must be tested or certified prior to use? At regular intervals thereafter?

- Are tapes and disks refinished or replaced before performance is degraded?
- Are air conditioners adequate for peak thermal loads? Are air conditioners backed up?
- Is there a schedule for frequent filter changes?
- Have all static electricity generators been disabled?
- Have all sources of water damage been eliminated?
- Is good housekeeping enforced throughout the facility?
- Is access to data terminals restricted?
- Are terminals and surrounding areas examined frequently to detect passwords carelessly left about?
- Is a log maintained of unsuccessful attempts to enter the computer from terminals?
- Is the log used to prevent further attempts?
- Is a log maintained of all successful entries to sensitive data?
- Is the log used to verify authorizations?
- Are terminals equipped with automatic identification generators?
- Are test procedures adequate to assure high-quality data transmissions?
- Is cryptography or scrambling used to protect sensitive data?
- Has a complete backup plan been formulated? Is it updated frequently?
- Does the backup plan include training, retraining, and cross-training of personnel?
- Is onsite backup available for the central processing unit? For peripherals?
- Does your backup site advise you of all changes to its hardware configuration and operating system?
- Does your backup site have enough free time available to accommodate your emergency needs?
- Do you monitor power-line voltage and frequency?
- Are the effects of brownouts, dim-outs, and blackouts known?
- Is advance warning available, and if so, is there a checklist of actions to be taken?
- Are power correctors in use? Voltage regulators? Line conditioners? Lightning spark gaps?
- Is backup power available? Dual substation supply? Motor generators? Uninterruptible power supplies?
- Does your equipment provide automatic restart and recovery after a power failure?
- Are backup plans tested realistically? At frequent intervals?

Microcomputers

In addition to the appropriate items just listed:

- Are removable disks always kept in a closed container when not actually mounted in a disk drive?
- Is it forbidden to put food or drink on or near computer equipment?
- Are personal computers securely fastened to prevent dropping or theft?
- Are air vents kept free?
- Are accurate inventories maintained?

- Is electrical power properly wired?
- Are uninterruptable power supplies in place?
- Has static electricity been eliminated?
- Are data communications secure?
- Is there an effective maintenance plan?

4.16 FURTHER READING

Ayers, J. E. *Digital Integrated Circuits: Analysis and Design*. Boca Raton, FL: CRC Press, 2003. (Second edition scheduled for publication in 2009.)

Clements, A. *Principles of Computer Hardware*, 4th ed. New York: Oxford University Press, 2006.

Horak, R. *Telecommunications and Data Communications Handbook*. Hoboken, NJ: Wiley-Interscience, 2007.

Kerns, D. V. *Essentials of Electrical and Computer Engineering*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 2004.

Pattern, D. A., and J. L. Hennessy. *Computer Organization and Design: The Hardware Software Interface*, 3rd ed. Los Angeles: Morgan Kaufmann, 2007.

Stallings, W. *Computer Organization and Architecture: Designing for Performance*, 7th ed. Upper Saddle River, NJ: Prentice-Hall, 2005.

CHAPTER 5

DATA COMMUNICATIONS AND INFORMATION SECURITY

Raymond Panko and Eric Fisher

5.1 INTRODUCTION 5·2

5.2 SAMPLING OF NETWORKS 5·2

5.2.1 Simple Home Network 5·2
5.2.2 Building LAN 5·4
5.2.3 Firms' Wide Area Networks (WANs) 5·5
5.2.4 Internet 5·7
5.2.5 Applications 5·9

5.3 NETWORK PROTOCOLS AND VULNERABILITIES 5·9

5.4 STANDARDS 5·9

5.4.1 Core Layers 5·10
5.4.2 Layered Standards Architectures 5·10
5.4.3 Single-Network Standards 5·11
5.4.4 Internetworking Standards 5·13

5.5 INTERNET PROTOCOL (IP) 5·14

5.5.1 IP Version 4 Packet 5·14
5.5.2 IP Version 6 5·16
5.5.3 IPsec 5·17

5.6 TRANSMISSION CONTROL PROTOCOL (TCP) 5·18

5.6.1 Connection-Oriented and Reliable Protocol 5·18
5.6.2 Reliability 5·20
5.6.3 Flag Fields 5·20
5.6.4 Octets and Sequence Number 5·21
5.6.5 Acknowledgment Numbers 5·21
5.6.6 Window Field 5·21
5.6.7 Options 5·21
5.6.8 Port Numbers 5·22
5.6.9 TCP Security 5·23

5.7 USER DATAGRAM PROTOCOL 5·23

5.8 TCP/IP SUPERVISORY STANDARDS 5·24

5.8.1 Internet Control Message Protocol (ICMP) 5·24
5.8.2 Domain Name System (DNS) 5·25
5.8.3 Dynamic Host Configuration Protocol (DHCP) 5·26
5.8.4 Dynamic Routing Protocols 5·27
5.8.5 Simple Network Management Protocol (SNMP) 5·27

5.9 APPLICATION STANDARDS 5·28

5.9.1 HTTP and HTML 5·28
5.9.2 E-Mail 5·28
5.9.3 Telnet, FTP, and SSH 5·28
5.9.4 Other Application Standards 5·29

5.10 CONCLUDING REMARKS 5·29

5.11 FURTHER READING 5·29

5.12 NOTES 5·29

5.1 INTRODUCTION. Sometimes, an attacker can simply walk up to a target computer. In most cases, however, attackers must use networks to reach their targets. Some attacks even aim *at* networks, trying to bring down local area networks, wide area networks, and even the global Internet. This chapter provides an overview of networking to help readers of this *Handbook* when they come across networking concepts in other chapters or in other contexts. This chapter covers a limited number of networking concepts. Specifically, it focuses on aspects of networking that are most relevant to security.

Before beginning, readers should note three important pieces of terminology that pervade the chapter.

1. This chapter often uses the term *octet*, which is a *byte*—a collection of eight bits. Networking grew out of electrical engineering, where octet is the preferred term; it is also widely used in the international technical community.
2. The second term is *host*. Any device attached to the global Internet is called a host. This includes everything from large server hosts to client PCs, personal digital assistants, mobile telephones, and even Internet-accessible coffeepots.
3. We will distinguish between the terms *internet* and *Internet*; the latter refers to the global Internet. However, *internet* spelled in lower case is either the Internet layer in the TCP architecture (see Section 5.6) or a collection of networks that is not the global Internet.

5.2 SAMPLING OF NETWORKS. This section looks briefly at a series of increasingly complex networks, giving the reader a high-level overview of what networks look like in the real world.

5.2.1 Simple Home Network. Exhibit 5.1 shows a simple home PC network. The home has two personal computers. The network allows the two PCs to share files and the family's single laser printer. The network also connects the two computers to the Internet.

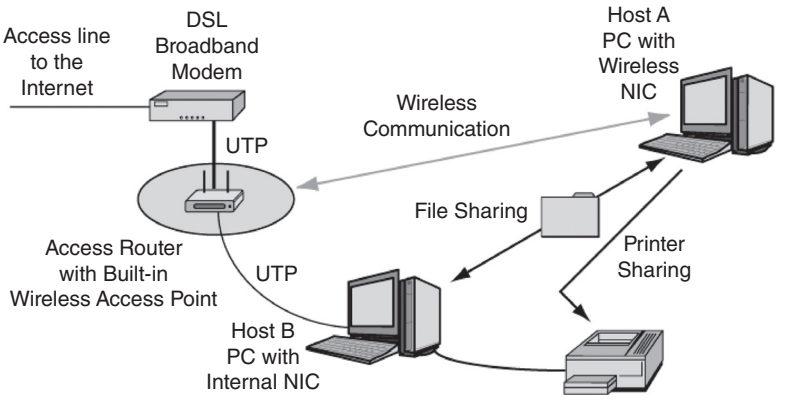


EXHIBIT 5.1 Simple Home Network

5.2.1.1 Access Router. The heart of this network is its *access router*. This small device performs a variety of functions, most importantly these five:

1. It performs as a switch. When one PC in the home sends messages (called packets) to the other hosts, the switch transfers the packets between them.
2. The access router is a wireless access point (WAP), which permits wireless computers to connect to it. Host A connects to the access router wirelessly.
3. A router connects a network to another network—in this case, it connects the internal network to the global Internet.
4. To use the Internet, each computer needs an Internet Protocol (IP) address. We will see later that IP is the main protocol that governs communication over the Internet. The access router has a built-in Dynamic Host Configuration Protocol (DHCP) server that gives each home PC an IP address.
5. The router provides *network address translation* (NAT), which hides internal IP addresses from potential attackers. Most routers also have a firewall for added security. WAPs are easily exploited if not configured with proper authentication security. Wireless signals can be transferred up to 800 feet away or more with special equipment. Without constant monitoring to defeat intrusions, an attacker can connect to an access point without the user's knowledge and intercept all passing traffic. Using NAT is essential to keeping a home network secure. Users should always enable this feature in order to prevent their hosts from being directly accessible to the public Internet, where direct scans and attacks are prevalent.

5.2.1.2 Personal Computers. Each of the two PCs needs circuitry to communicate over the network. Traditionally, this circuitry came in the form of a printed circuit board, so the circuitry was called the computer's *network interface card* (NIC). In most computers today, the circuitry is built into the computer; there is no separate printed circuit board. However, the circuitry is still called the computer's NIC.

In this small network, the two computers share their files. Given the wireless access capability of the network, drive-by hackers could potentially read shared files as well. File sharing without strong wireless security is dangerous. It is important to set up *Wi-Fi Protected Access* (WPA or WPA2) or 802.11i security in pre-shared key (PSK) mode on both the access router/access point and each of the client PCs.

It is important to configure the PCs for security. Although NAT by itself is strong, and most routers also provide stateful-inspection firewalls (see Chapter 26 in this *Handbook*), some attacks will inevitably get through to the internal network. Hosts must have strong firewalls, antivirus programs, and antispyware programs (see Chapter 41); and they must be updated automatically when security patches are released by the operating system vendor and by application program vendors (see Chapter 40).

5.2.1.3 UTP Wiring. In Exhibit 5.1, Host B connects to the access router via copper wiring called a UTP cable, Ethernet (IEEE 802.3) cable, or commonly Cat5 (Cat5 stands for Category 5 cabling, defined in standard ANSI/TIA/EIA-568-A). It uses four-pair *unshielded twisted pair* (UTP) wiring inside the cord jacket. As Exhibit 5.2 shows, a UTP cord contains eight copper wires organized as four pairs. The two wires of each pair are twisted around each other. The *RJ-45* connectors at the ends of a UTP cord look like RJ-11 home telephone connectors but are a little wider. (*RJ* means

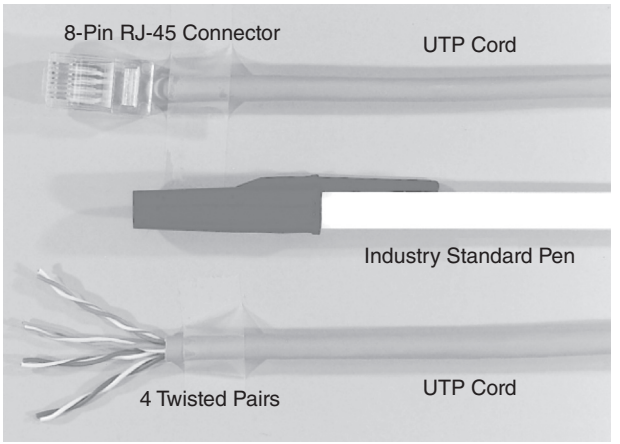


EXHIBIT 5.2 Unshielded Twisted Pair (UTP) Wiring Cord

Registered Jack and originally referred to Bell System order codes; it is now defined by the Administrative Council for Terminal Attachment, ACTA.)

5.2.1.4 Internet Access Line. The home network needs an Internet access line to connect the home to the Internet. In Exhibit 5.1, this access line is a *digital subscriber line* (DSL) high-speed access line, and the home connects to this access line via a small box called a *DSL modem*. (The DSL modem connects to the access router via a UTP cord; it connects to the wall jack via an ordinary telephone cord.) Other Internet access technologies include slow telephone modems, fast cable modems, geosynchronous-satellite connections, and even wireless access systems. Most of these technologies are called *broadband access lines*. In general, *broadband* simply means *very fast*, although in radio transmission it describes a wide range of frequencies.

5.2.2 Building LAN. The home network shown in Exhibit 5.1 is a *local area network* (LAN). A LAN operates on a customer's premises—the property owned by the LAN user. (For historical reasons, *premises* is always spelled in the plural.) In the case of the home network, the premises consist of the user's home or apartment. Exhibit 5.3 shows a much larger LAN. Here, the premises consist of a corporate multistory office building.

On each floor, computers connect to the floor's workgroup switch via a UTP cord or a wireless access point. The workgroup switch on each floor connects to a *core switch* in the basement equipment room. The router in the basement connects the building LAN to the outside world.

Suppose that Client A on Floor 1 sends a packet to Server X on Floor 2. Client A sends the packet to Workgroup Switch 1 on the first floor. That workgroup switch sends the packet down to the core switch in the basement. The core switch then sends the packet up to Workgroup Switch 2, which passes the packet to Server X.

UTP is easy to wiretap, allowing attackers to read all packets flowing through the cord. Telecommunications closets should be kept locked at all times, and cords should be run through thick metal wiring conduits wherever possible (for more details of physical and facilities security, see Chapters 22 and 23 in this *Handbook*). UTP also generates weak radio signals when traffic flows through it. It is possible to read these signals from some distance away using highly specialized equipment. Newer

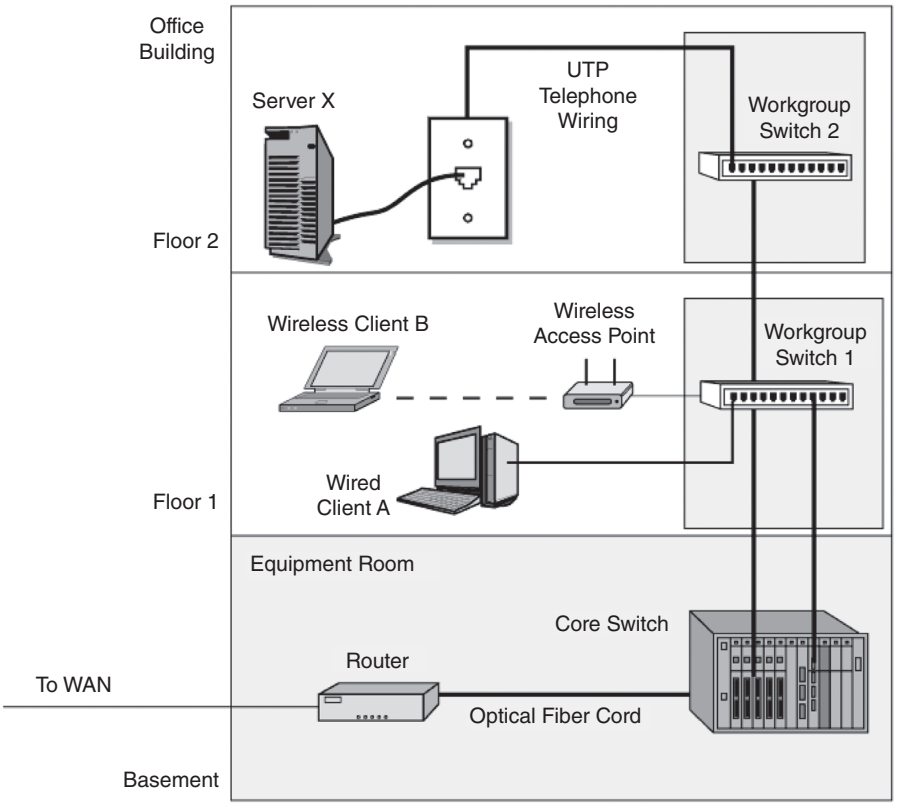


EXHIBIT 5.3 Building LAN

specifications called Cat5e and Cat6 were developed to cut down on interference and cables can be purchased with shielding, but even then it is possible to eavesdrop.

Eavesdropping by tapping a UTP cable is not difficult once physical access is gained; however, typically there are far easier ways of gaining access to a network and far more desirable targets. Eavesdropping on a wire would reveal any passing traffic, but eavesdropping on a router or switch would reveal passing traffic on *many* wires. Physical security is an important facet of network security and must be properly addressed, but most attacks today rely on more virtual vulnerabilities.

For more extensive details of LAN security, see Chapter 25 in this *Handbook*.

5.2.3 Firms' Wide Area Networks (WANs). Although LANs operate within a company's premises, *wide area networks* (WANs) connect geographically separate sites—usually within a single corporation. Corporations do not have the regulatory rights-of-way needed to run wires though public areas. For WAN service, companies must use companies called *carriers* that do have these rights-of-way.

Exhibit 5.4 shows that most firms use multiple-carrier WANs. In the exhibit, some sites in this company are connected by point-to-point *leased lines* from a telephone company. The companies also subscribe to *switched network services* that transfer traffic between several sites. The exhibit shows that these switched network services use the *Frame Relay* technology. The company uses two separate Frame Relay networks—one to connect its own sites to one another and another to connect it to another firm.

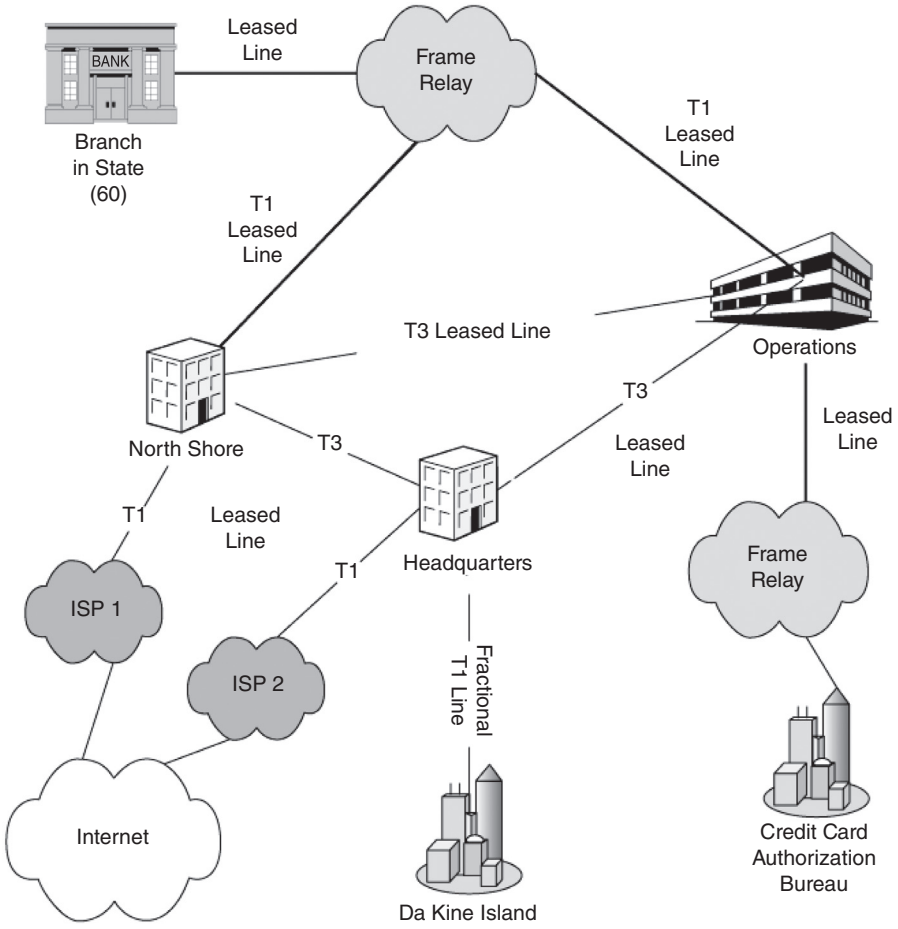


Exhibit 5.4 Wide Area Networks (WANs)

Carrier technology is usually considered more secure by security professionals due to its closed-access nature. Unlike the Internet, which allows anyone to connect to it, only commercial firms may connect to carrier WANs, which makes attacker access very difficult. However, attacker access is not impossible. For example, if an attacker hacks a computer owned by the carrier (or even by a customer), this breach may permit access.

In addition, the carrier alone knows how it routes traffic through its network. This should stymie attackers even if they somehow get access to the network. However, such *security through obscurity* is considered a poor practice by security professionals because it is possible for attackers who hack carrier computers to get access to routing information. (Attackers usually have much simpler attack vectors; see Chapters 15 and 19 in this *Handbook* for more details.)

Although carrier technology is more secure, it is also extremely expensive. With the development of virtual private networks (VPNs), companies can connect geographically disparate groups of computers *virtually* over the common internet. This provides much of the security benefit of WANs while cutting the implementation cost dramatically. See Chapter 32 for more information about VPN security.

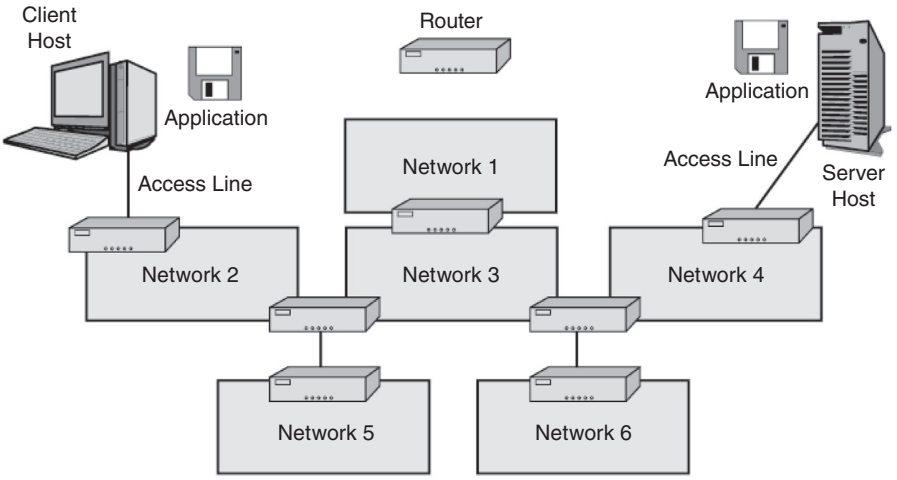


EXHIBIT 5.5 Internet

5.2.4 Internet. By the end of the 1970s, there were many LANs and WANs in the world. Many of the WANs were nonprofit networks that connected universities and research institutions. Unfortunately, computers on one network could not talk to computers on other networks. To address this problem, the Defense Advanced Research Projects Agency (DARPA) created ARPANET in 1969, the origin of today's Internet, based on the pioneering conceptual design for what J. C. R. Licklider called the *Intergalactic Computer Network* in a 1963 paper. By definition, *an internet* connects individual networks together. Later, commercial networks were allowed to join later versions of ARPANET, and it became the Internet we know today.

Exhibit 5.5 shows that devices called *routers* connect the individual networks together. Initially, these devices were called *gateways*. The term *gateway* was used instead of "router" in some early standards, but most vendors have now adopted the name "router." There are two exceptions, the first being Microsoft, which still tends to call routers "gateways." The second is the router directly accessible to a network, and thus the first hop when exiting a network is often called the *default gateway*.

Any computer on any network on the Internet can send messages to any computer on any other network on the Internet. The messages that travel all the way from one computer to another across the Internet are called *packets*.

Exhibit 5.6 shows that the packet travels all the way from the source host to the destination host. Along the way, it is routed through different networks until arriving at its destination.

The global Internet uses a suite of communication protocols known as *Transmission Control Protocol/Internet Protocol* (TCP/IP). In addition, many firms build separate internal TCP/IP networks for their own communication. These internal networks are called *intra*nets to distinguish them from the *Inter*net.

Initially, security on internal networks was comparatively light because it was assumed that external attackers would have a difficult time getting into corporate intranets. However, if a hacker takes over an internal computer connected to the intranet, light security becomes a serious problem. Consequently, most firms have been progressively hardening their intranet security.

Exhibit 5.7 shows that individual homes and corporations connect to the Internet via carriers called *Internet service providers* (ISPs). The Internet has many ISPs, but

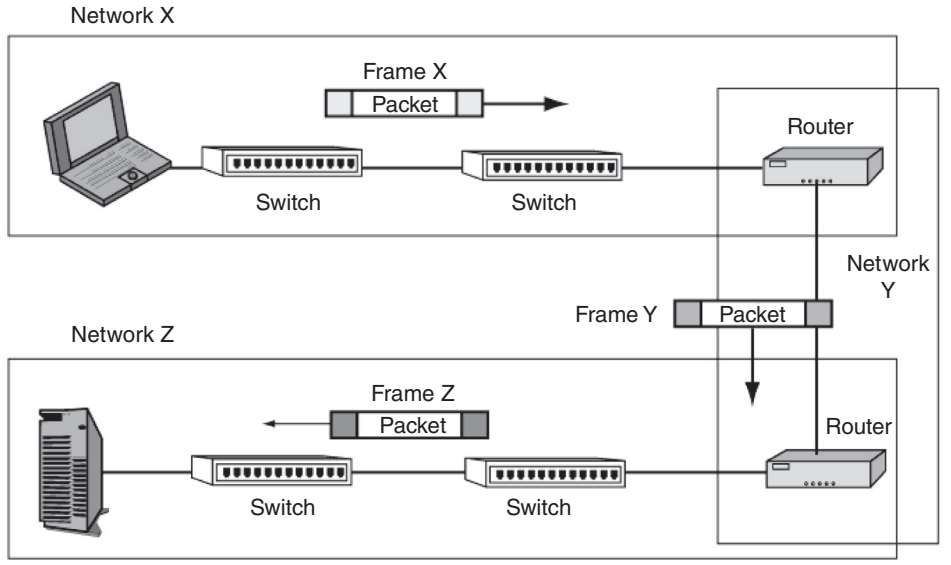


EXHIBIT 5.6 Frames and Packets

they all connect at centers that usually are called *network access points* (NAPs). These connections allow global communications for all connected hosts.

Most ISPs are commercial organizations run for profit to provide Internet access for home users. There is no central access control to the Internet; however, there are central agencies for controlling Domain Name Systems (DNSs) called *registrars*.

When the Internet was designed in the late 1970s, there was a conscious decision to promote openness and not to add the burdens of security. As a consequence of a lack of security technology and open access to almost anyone, the Internet is a security nightmare. Companies that transmit sensitive information over the Internet need to consider cryptographic protections. (See Chapters 7, 32, 33, 34, 35, and 37 in this *Handbook* for more details of cryptography and other means for achieving security on networks.)

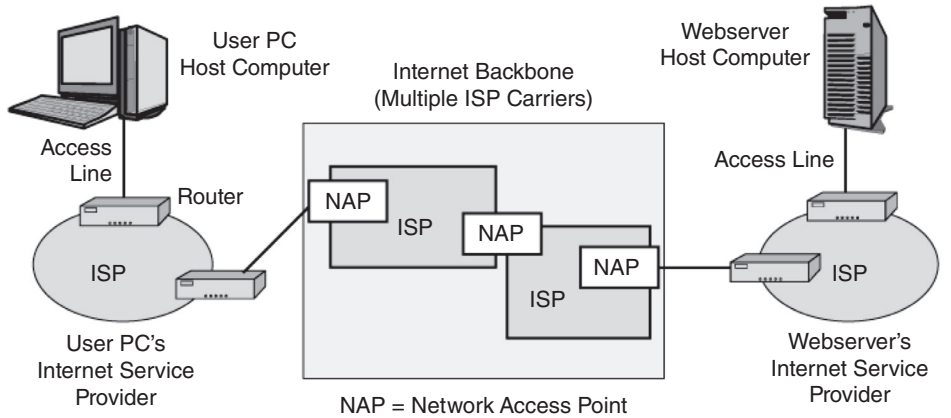


EXHIBIT 5.7 Internet Service Providers (ISPs)

5.2.5 Applications. Although the inner workings of the Internet rely on networks, most users are only aware of the applications they commonly use that run on top of networks. Familiar personal applications include the World Wide Web, email, and instant messaging, among many others. Corporations use some of these applications, but they also use many business-specific applications, such as accounting, payroll, billing, and inventory management. Often, business applications are transaction-processing applications, which are characterized by high volumes of simple repetitive transactions. The traffic volume generated by transaction-processing and other business-oriented applications usually far outweighs the traffic of personal applications in the firm. (See Chapter 30 in this *Handbook* for details of e-commerce security.)

All programs have bugs, including security vulnerabilities. There are many applications, and keeping track of application vulnerabilities and constantly patching many applications is an enormous task that is all too easy to put off or complete only partially. (See Chapter 40 for an overview of patch management.) Also, each application must be configured with options that have high security, and security must be managed on each application (e.g., anti-virus and spam blocking in email). (See Chapter 20 for a review of spam and anti-spam measures.)

5.3 NETWORK PROTOCOLS AND VULNERABILITIES. The products of different network vendors must be able to work together (interoperate). This is possible only if there are strong communication standards to govern how hardware and software processes interact. With such standards, two or more programs can interoperate effectively.

Standards raise three security issues. One is the standard itself. For instance, the TCP standard discussed later in this chapter is difficult to attack because an attacker cannot send a false message unless he or she can guess the sequence number of the next message. This normally is very difficult to do. However, if the attacker sends an RST (reset) message, which terminates a connection, this protection is greatly reduced. In fact, it is fairly easy to send RST messages that close legitimate open connections.

A second issue is security built into the standard. Most standards were created without security, and security was added only in later versions, sometimes in an awkward way. For instance, IP, which is the main protocol for delivering packets over the Internet, originally had no security. The *IP security* (*IPsec,* pronounced *eye-pea-sek*) standards were created to address this weakness, but IPsec is burdensome and not widely used.

Another security weakness of early versions of IP, including the widely used IPv4, is the limitation on address space due to the 32 bit address field in the IPv4 packet (yielding an address space of about 4×10^9); as this edition of this *Handbook* goes to press, IPv4 address exhaustion is being addressed by the migration to IPv6, with its 128-bit addresses (an address space of about 3×10^{38}).

A further issue is the security of the implementation of standards in vendor products. Most attacks that aim at standards weaknesses attack vendor products that have security vulnerabilities unrelated to the protocols they implement.

5.4 STANDARDS. Networks and network security depend on standards. Standards have permitted global interconnectivity, unlike the early years of networking when proprietary products dominated the world of computing and interconnection was difficult or impossible.

EXHIBIT 5.8 Three Standards Core Layers

Super Layer	Description
Application	Communication between application programs on different hosts attached to different networks on a network.
Internetworking	Transmission of packets across a routed internet. Packets contain application-layer messages.
Single network	Transmission of packets across a single-switched network.

5.4.1 Core Layers. Standards are complex, and when people deal with complex problems, they usually break these problems into smaller parts and have different specialists work on the different parts. Exhibit 5.8 shows that standards are divided into three core layers that collectively have the functionality needed to allow an application program on one network in an internet to interoperate with another program on another computer on another network.

At the *application core layer*, the two applications must be able to interact effectively. For instance, in World Wide Web access, the two application programs are the browser on the client PC and the Web server program on the Web server. The standard for Web interactions is the *Hypertext Transfer Protocol* (HTTP). Both the browser and the Web server applications have to send messages that comply with the HTTP standard.

The middle layer is the *internet core layer*. Standards at this layer govern how packets are delivered across a routed internet. One of the main standards at the internet core layer is the Internet Protocol (IP). We will see other internetworking standards later.

The lowest core layer is the *single-network core layer*. Standards at this layer govern the transmission of packets across the switches and transmission lines in a single-switched network (a LAN or WAN).

5.4.2 Layered Standards Architectures. Standards are created by standards agencies. These standards agencies first create detailed layering plans for creating standards. These specific layering plans are called *layered standards architectures*. Afterward, standards agencies create standards in the individual layers. Exhibit 5.9 shows two popular layered standards architectures and relates these standards architectures to the three core layers we saw earlier.

The *Internet Engineering Task Force* (IETF) is the standards agency for the Internet. Its standards architecture is called TCP/IP—a name taken from two of its most

EXHIBIT 5.9 Layered Standards Architectures

Super Layer	TCP/IP	OSI	Hybrid TCP/IP-OSI
Application	Application	Application Presentation Session	Application
Internet	Transport Internet	Transport Network	Transport Internet
Network	Subnet access	Data link Physical	Data link Physical

important standards, TCP and IP. Exhibit 5.9 shows that TCP/IP has four layers. The bottom layer, the *subnet access layer*, corresponds to the single-network core layer. The top layer, in turn, is the *application layer*, which corresponds to the application core layer. The two middle layers—the *internet* and *transport* layers—correspond to the internet core layer. TCP/IP focuses primarily on internet working. Dividing this core layer into two TCP/IP layers permits greater division of labor in standards development.

The other standards architecture shown in the figure is *OSI*, which is rarely spelled out by its full name, the *Reference Model of Open Systems Interconnection*. OSI is governed by two standards agencies. One is ISO, the *International Organization for Standardization*. The other is ITU-T, the *International Telecommunications Union–Telecommunications Standards Sector*. (The official names and the official acronyms do not match because they originated in different languages.)

Exhibit 5.9 shows that OSI divides the three core layers into a total of seven layers. OSI single networks use standards at two layers—the *physical* and *data link* layers. OSI's market dominance is so strong at the physical and data link layers that the IETF rarely develops standards at these layers. The *subnet access* indication in the TCP/IP framework basically means *Use OSI standards here*.

Neither of these two standards architectures dominates. What nearly all firms use today is the hybrid TCP/IP–OSI standards architecture, which Exhibit 5.9 illustrates. This hybrid architecture uses OSI standards at the physical and data link layer and TCP/IP standards at the internet and transport layer. Corporations also use standards from some other standards architectures at the internet and transport layers, but TCP/IP standards dominate.

At the application core layer, the situation is complex. Both OSI and TCP/IP standards are used, often in combination. In fact, OSI standards often reference TCP/IP standards and vice versa. Although OSI and TCP/IP are often viewed as rivals, this is not the case at all. Several other standards agencies also create application layer standards, complicating the picture even further.

5.4.3 Single-Network Standards. As just noted, OSI standards dominate in the two single-network layers—the physical and data link layers. Exhibit 5.10 shows how the physical and data link layers are related.

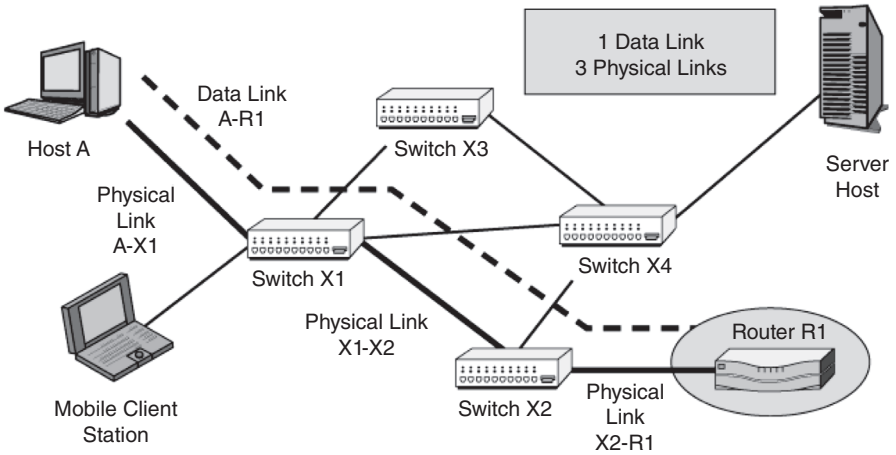


EXHIBIT 5.10 Physical and Data Link Layers

5.4.3.1 Data Link Layer. The path that a frame takes through a single network is called the frame's data link. In Exhibit 5.10, the data link runs between Host A and Router R1. This data link passes through Switch X1 and Switch X2.

The source computer sends the frame to the first switch, which forwards the frame to the next switch along the data link, which forwards the frame further. The last switch along the data link passes the frame to the destination computer (or router, if the packet in the frame is destined for a computer on another network).

5.4.3.2 Physical Layer. Physical-layer standards govern the physical connections between consecutive devices along a data link. In Exhibit 5.10, these physical links are A–X1, X1–X2, and X2–R1. Earlier, we saw one popular transmission medium, unshielded twisted pair wire in Cat5 cables. UTP dominates in links between computers and workgroup switches (see Exhibit 5.3). UTP signals typically involve voltage changes. For instance, a high voltage may indicate a 1, while a low voltage may indicate a 0. (Actual voltage patterns usually are much more complex.)

For longer distances and very high speeds, another popular transmission medium is optical fiber, which sends light signals through thin glass tubes. Optical fiber signals actually are very simple. In a clock cycle, the light is turned on for a 1 or off for a 0.

UTP cords act like radio antennas when they carry signals. Some of the signal always radiates out, allowing people to intercept transmission signals by placing devices near (but not touching) the cord. Intercepting and interpreting electromagnetic emissions from computing devices is called van Eck phreaking (also famously codenamed "TEMPEST" by the NSA) after the Dutch scientist Wim van Eck published a paper in 1985 demonstrating how to monitor and reconstitute leaked signals from *cathode-ray terminals* (CRTs). In contrast, optical fiber requires physically tapping into the fiber cords. Physical wiretapping can also be done with UTP, but there are often far easier methods to intercept or steal traffic rather than trying to physically tap the wires.

Wireless transmission uses radio waves. This permits mobile devices to be served in ways never before possible. Wireless transmission is used for both LAN and WAN transmission.

Radio signals spread widely, even when dish antennas are used. Consequently, it is very easy for eavesdroppers to listen in on radio transmissions and do other mischief. Radio signals must be strongly encrypted, and the parties must be strongly authenticated to prevent impostors from sending radio transmission.

Radio signaling is very complex. Most radio signaling uses spread spectrum transmission, in which the information is sent over a wide range of frequencies. *Spread spectrum* transmission is used to improve propagation reliability. Radio transmission has many propagation problems, such as interference from other sources. Many propagation problems occur only at certain frequencies. By spreading the signal across a wide spectrum of frequencies and doing so redundantly, the signal will still be intelligible even if there are strong problems at some frequencies.

Prabakar Prabakaran summarized the benefits of spread-spectrum communications as follows:

> Spread-spectrum systems provide some clear advantages to designers … [H]ere are nine benefits that designers can expect when using a spread-spectrum-based wireless system.
>
> **1.** Reduced crosstalk interference: In spread-spectrum systems, crosstalk interference is greatly attenuated due to the processing gain of the spread spectrum system as described earlier …
> **2.** Better voice quality/data integrity and less static noise …

3. Lowered susceptibility to multipath fading …
4. Inherent security: In a spread spectrum system, a PN [pseudo-random number] sequence is used to either modulate the signal in the time domain (direct sequence systems) or select the carrier frequency (frequency hopping systems). Due to the pseudo-random nature of the PN sequence, the signal in the air has been "randomized." Only the receiver having the exact same pseudo-random sequence and synchronous timing can de-spread and retrieve the original signal. Consequently, a spread spectrum system provides signal security that is not available to conventional analog wireless systems.
5. Co-existence: A spread spectrum system is less susceptible to interference than other non-spread spectrum systems. In addition, with the proper designing of pseudo-random sequences, multiple spread spectrum systems can co-exist without creating severe interference to other systems. This further increases the system capacity for spread spectrum systems or devices.
6. Longer operating distances …
7. Hard to detect: Spread-spectrum signals are much wider than conventional narrowband transmission (of the order of 20 to 254 times the bandwidth of narrowband transmissions). Since the communication band is spread, it can be transmitted at a low power without being detrimentally by background noise …
8. Hard to intercept or demodulate: The very foundation of the spreading technique is the code used to spread the signal …
9. Harder to jam: The most important feature of spread spectrum is its ability to reject interference … .[1]

The military uses frequency-hopping spread-spectrum (FHSS) transmission for security. Military spread-spectrum transmission works in such a way that makes intercepting transmissions very difficult. Civilian spread-spectrum transmission, in contrast, is designed to make connecting simple and therefore offers relatively little security.

Switches spend almost all of their time forwarding frames. However, switches spend some of their time exchanging supervisory information packets with one another to keep the network running efficiently. For example, in *Ethernet* (IEEE 802.3), which dominates LAN standards, if there are loops among the switches, the network will malfunction. If a switch detects a loop, it sends supervisory packets to other switches. The switches in the network then communicate until they determine the most appropriate path and disable other ports to prevent the internal looping. This process is governed by the *Spanning Tree Protocol* (STP, part of IEEE 802.1) or the newer *Rapid Spanning Tree Protocol* (RSTP, defined in IEEE 802.1 w and now part of IEEE 802.1D-2004).

Attackers can create denial-of-service (DOS) attacks on the switches in a network by impersonating a switch and sending a flood of false messages to the network's real switches indicating the presence of a loop. The switches may spend so much of their time reorganizing the network that they will be unable to serve legitimate traffic. They also can attack several other supervisory protocols to make switches unavailable for processing normal packets. The 802.1AE standard is designed to limit switch-to-switch communication to authenticated switches.

5.4.4 Internetworking Standards. As noted earlier, the IETF divided the internetworking core layer into two layers—the internet and transport layers. Exhibit 5.11 shows how the two layers are related.

The internet layer forwards packets, hop by hop, among routers until the packet reaches the destination host. The main standard at the internet layer is the Internet Protocol (IP).

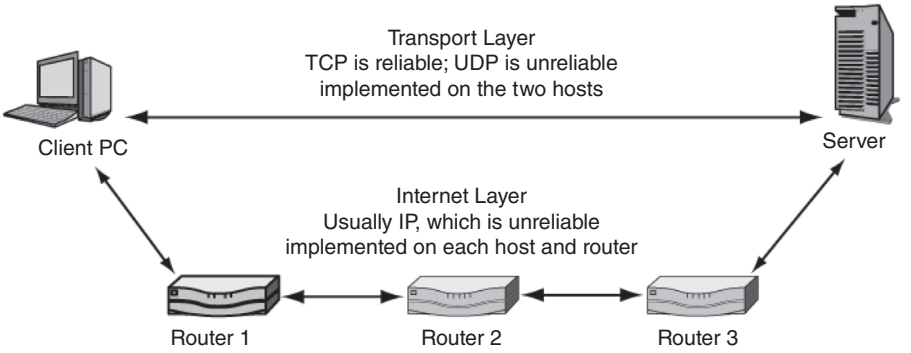


EXHIBIT 5.11 Internet- and Transport-Layer Standards

The designers of TCP/IP realized that they could not predict what services the single networks connecting routers would provide. IP was made a simple best-effort protocol, in order to assume minimal functionality in the single networks along the way. There are no guarantees that packets will arrive at all or, if they do arrive, that they will arrive in order.

To make up for the limitations of IP, a transport layer was added. The main standard designed for this layer, the *Transmission Control Protocol* (TCP), was created as a high-capability protocol that would fix any errors made along the way, ensure that packets arrived in order, slow transmission when the network became overloaded, and do several other things. For applications that did not need this level of reliability, a simpler standard was created, the *User Datagram Protocol* (UDP).

5.5 INTERNET PROTOCOL (IP). The Internet Protocol (IP) does two main things. First, it governs how packets are organized. Second, it determines how routers along the way move packets to the destination host. (Analogously, data-link-layer standards govern how frames containing packets are organized and how switches along the way move the frame across a single-switched network.)

5.5.1 IP Version 4 Packet. The main version of the Internet Protocol is Version 4 (IPv4). (There were no Versions 0, 1, 2, or 3.) This version has been in use since its definition in 1981 and will continue to be used for many years to come, although IPv6 is intended to supersede it. Exhibit 5.12 shows the IPv4 packet's organization.

A packet is a long stream of 1s and 0s. The IP *header* normally is shown on several rows, with 32 bits on each row. The first row has bits 0 through 31; the next row shows bits 32 through 63 and so on.

The header is divided into smaller units called *fields*. Fields are defined by their bit position in the packet. For example, the first four bits comprise the *version number* field. These are bits 0 through 3. In IPv4, this field holds 0100, which is 4 in binary. The *header length* field comprises the next four bits (bits 3 through 7).

5.5.1.1 First Row. As just noted, the first field (bits 0 through 3) is the version number field. In IPv4, the value is 0100 (4). In the newer version of the Internet Protocol, *IP Version 6* (IPv6), the value is 0110.

The next field is the header length field. This gives the length of the headers in 32-bit units. As Exhibit 5.12 shows, a header without options has five 32-bit lines, so this field will have the value 0101 (5 in binary).

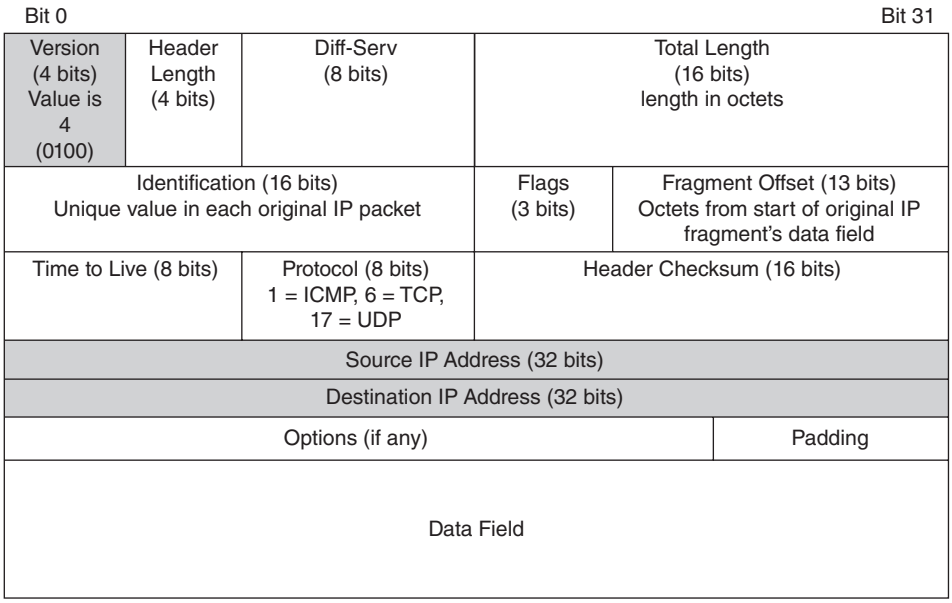


EXHIBIT 5.12 Internet Protocol (IP) Packet

The use of options is uncommon in practice. In fact, options tend to indicate attacks. Therefore, a value larger than 5 in the header length field indicates that the packet header has options and is therefore suspicious.

The 1-octet *dif-serv* (differential services) field was created to allow different services (priority, etc.) to be given to this packet. However, this field typically is not used.

The *total length* field gives the length of the entire IP packet in octets (bytes). Given the 16-bit length of this field, the maximum number of octets in the IP packet is 65,536 (216). Most IP packets, however, are far smaller. The length of the data field is this total length minus the length of the header in octets.

5.5.1.2 Second Row. If an IP packet is too long for a single network along the way, the router sending the packet into that network will fragment the packet, dividing its contents into a number of smaller packets. For assembly on the destination host, all fragment packets are given the same *identification field* value as in the original packet. The data octets in the original packets are numbered, and the number of the first data octet in the packet is given a *fragment offset* value (13 bits long). There are three *flag* fields (1-bit fields). One of these, *more fragments*, is set to 1 in all but the last packet, in which it is made 0. The information in these three fields allows the destination host to place the packets in order and know when there are no more packets to arrive.

IP fragmentation by routers is usually rare, and attackers can use fragmentation to hide attack information. Even if the first fragment packet is dropped by the firewall, other packets that do not have the signature information in the first header can get through. Therefore, IP fragmentation is suspicious.

5.5.1.3 Third Row. The third line begins with an ominous-sounding *time to live* (TTL) field, which has a value between 0 and 255. The sending host sets the initial value (64 or 128 in most operating systems). Each router along the way decreases the

value by 1. If a router decreases the value to 0, it discards the packet. This process was created to prevent misaddressed packets from circulating endlessly around the Internet.

To identify hosts, attackers will use the Internet Control Message Protocol (ICMP) to *ping* many IP addresses (as discussed in Section 5.8.1). A reply tells the attacker that a host exists with that IP address. In addition, by guessing the initial TTL value and looking at the TTL value in the arriving packet, the attacker can guess how many router hops separate the attacker's host from the victim host. Sending many pings to different IP addresses can help the attacker map the routers in the target network. Often, administrators will turn off ICMP traffic outside of the internal networks in order to prevent anyone who isn't an authorized user from mapping active internal hosts.

The *data* field of the IP packet may contain a TCP message segment, a UDP datagram message, or something else, such as the ICMP messages we will discuss in Section 5.8.1. A value of 1 in this field indicates that the data field is an ICMP message. In turn, 6 indicates a TCP segment, and 17 indicates that the data field contains a UDP header.

The *header checksum* field contains a value placed there by the sender. This number is determined by a calculation based on the values of other fields. The receiving internet process redoes the calculation. If the two numbers are different, then there must have been an error along the way. If so, the router or destination host receiving the packet will simply discard the packet. There is no retransmission, so *IP* is not inherently reliable; however, one of the functions of *TCP* is to monitor the sequence numbers and initiate retransmission of missing packets. See Section 5.6.2.

5.5.1.4 Source and Destination IP Address. When you send a letter, the envelope has an address and a return address. The analogous addresses in IP headers are the source and destination IP addresses. Note that IP addresses are 32 bits long. For human reading, these 32 bits are divided into four 8-bit *segments*, and each segment's bits are converted into a decimal number between 0 and 255. The four segment numbers are then separated by dots. An example is *128.171.17.13*. Note that this dotted decimal notation is a memory and writing crutch for inferior biological entities (people). Computers and routers work with 32-bit IP addresses directly.

Many forms of firewall filtering are based on IP addresses. In addition, many attackers spoof their packet's source IP address (i.e., replace the real IP address with a false IP address).

5.5.2 IP Version 6. Although IP Version 4 is widely used, its 32-bit IP address size causes problems: It can address only 4,294,967,296 ($\sim 10^9$) devices. This relatively small size limits the number of possible IP addresses. In addition, when IP addresses were distributed, most addresses were assigned to the United States because the Internet was invented there. In fact, some U.S. universities received more IP addresses than China.

To address the limitations of the 32-bit IP address size, a new version of the Internet Protocol was created. This is *IP Version 6* (IPv6). (A Version 5 was defined, but it was never used.) Exhibit 5.13 shows the IPv6 packet organization.

One obvious change is that the IP addresses are much larger—128 bits. Each IP address, then, requires four 32-bit lines to write and is equivalent to $\sim 10^{38}$. This will provide IP addresses to allow almost every device to be a host on the Internet—including toasters and coffeepots. To give us a sense of the scale of this enormous number, it is enough to address every single molecule of water in a cube over 2 km on a side. Another popular description of the difference in size of the IPv4 and IPv6 address space is that if the address space of IPv4 were represented as a square roughly 4 cm on a side,

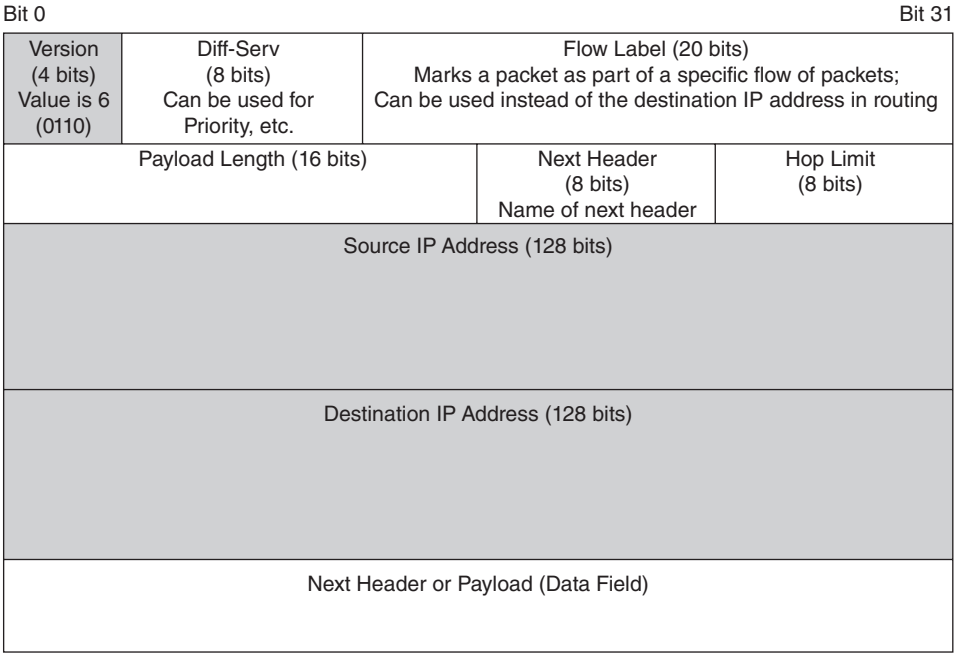


Exhibit 5.13 IP Version 6 Packet

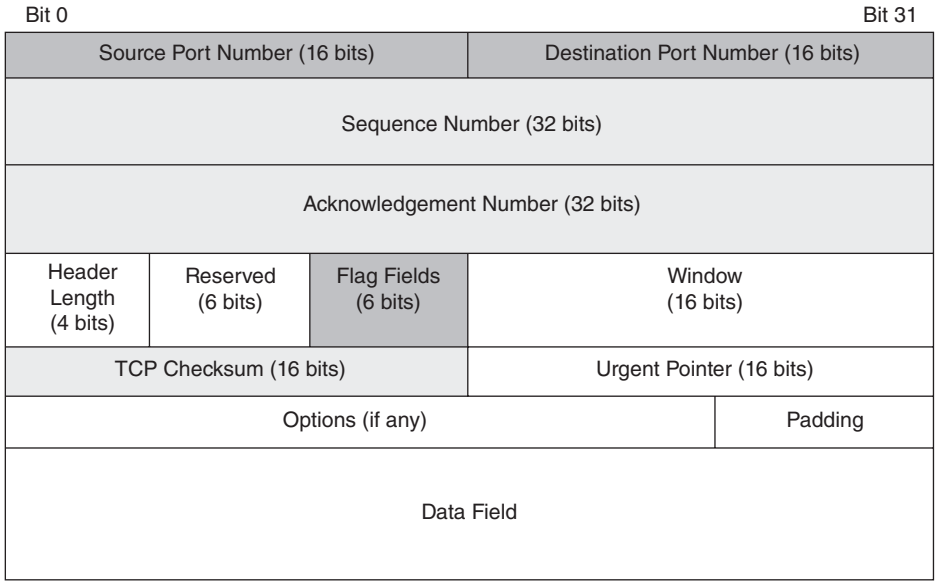
the equivalent area for IPv6 address space would cover the solar system out to Pluto's orbit.

The version number field is 4 bits long, and its value is 6 (0110). There also is a *dif-serv* field and a *flow label* field that is 20 bits long. These fields allow the packet to be assigned to a category of packets with similar needs. All packets in this category would be assigned the same flow label and would be treated the same way by routers. However, this capability is not widely used.

There is a *hop limit* field that serves the same function as the time to live (TTL) field in IPv4. The *payload length*, in turn, gives the length of the data field in octets.

A major innovation in IPv6 is the *next header* field. There can be multiple headers following the first header shown in Exhibit 5.13. For instance, IPsec security is implemented with a security header. Although options are unusual in IPv4, IPv6 uses additional headers extensively. The next header field tells what the next header is. Each additional header has a *next header* field that identifies the next header or says that there is no next header.

5.5.3 IPsec. IP, which was created in the early 1980s, initially had no security at all. Finally, in the 1990s, the Internet Engineering Task Force developed a general way to secure IP transmission. This was IP security, which normally is just called IPsec. IPsec functions by protecting a packet or most of a packet and sending the protected packet inside another packet. IPsec is a general security solution because everything within the data field of the protected packet is securely encrypted, including the transport and application layer information. This includes the transport message and the application message contained in the transport message. Originally developed for IPv6, it was extended to IPv4 as well, becoming a completely general solution. See Chapter 32 in this *Handbook* for further discussion of IPsec.



Flag fields are 1-bit fields. They include SYN, ACK, FIN, RST, PSH, and URG.

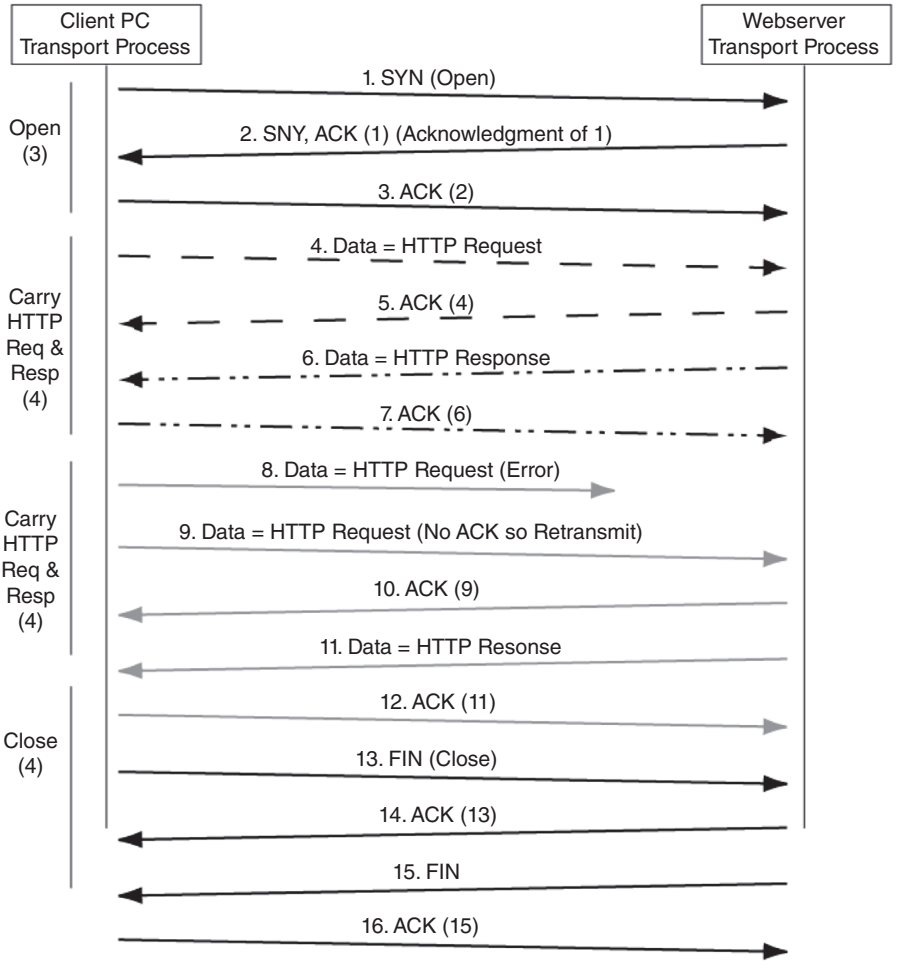
EXHIBIT 5.14 Transmission Control Protocol (TCP) Segment

5.6 TRANSMISSION CONTROL PROTOCOL (TCP). As noted earlier, the Transmission Control Protocol (TCP) is one of the two possible TCP/IP protocols at the transport layer. Exhibit 5.14 shows the TCP message, which is called a TCP *segment*.

5.6.1 Connection-Oriented and Reliable Protocol. Protocols are either connectionless or connection-oriented.

- Connection-oriented protocols are like telephone conversations. When you call someone, there is at least tacit agreement at the beginning of the conversation that you are able to speak. Explicit indicators such as "Hold, please." and "Can I call you back?" indicate an unwillingness to proceed at the moment. Also, there is at least tacit agreement that you are done talking at the end of the conversation; simply hanging up is considered rude. "Bye" or "Talk to you later" are examples of termination signals.
- Connectionless protocols, in turn, are like email. When you send a message, there is no prior agreement, and after the message is sent, there is no built-in provision for a reply (unless you are one of those people who asks to be notified when the receiver reads the message).

Exhibit 5.15 shows a sample TCP connection. Three messages are sent to open a connection. The originator sends a TCP SYN segment to indicate that it wishes to open a TCP session. The other transport process sends back a TCP SYN/ACK segment that acknowledges the connection opening message and indicates that it is willing to open the connection. The originator then sends an ACK segment to indicate reception of the SYN/ACK segment.



Note: An ACK may be combined with the next message if the next message is sent quickly enough.

EXHIBIT 5.15 Messages in a TCP Session

Attackers can use TCP connection openings to execute denial-of-service attacks that make a server unable to respond to legitimate traffic. The attacker sends a *SYN* segment to open a connection to the victim server. The victim server responds with a *SYN/ACK* message. The victim server also sets aside resources for the connection. The attacker never responds with an *ACK*, so this is called a *half-open SYN attack*. If the attacker floods a server host with SYN segments, the victim server will reserve so many resources that it will be overloaded and unable to serve legitimate connection opening attempts. The server may even crash. See Chapter 18 for discussion of denial-of-service attacks.

Ending a conversation, in contrast, normally takes four messages. One side sends a *FIN* segment, which the other party acknowledges. Then the other party sends a FIN segment, which the other side acknowledges. After the first side sends the original FIN segment, it will not send any new information, but it will send acknowledgments for segments sent by the other party.

There is another way to end a session or even to reject opening one. At any point, either party can send a *RST* (reset) message. An RST message ends the conversation abruptly. There is not even an acknowledgment. It is like hanging up in a telephone conversation.

Attackers often preface an attack by attempting to identify the IP addresses of running hosts—much like thieves casing a neighborhood. One way to do this is to send TCP SYN segments to hosts. If hosts reject the SYN segment, they often send back an RST message. As noted earlier, TCP segments are carried in the data fields of IP packets. The source IP address in the packet delivering the TCP RST segment will be that of the internal host. Whenever the attacker receives an RST segment, this verifies the existence of a working host at that packet's IP address. Firewalls often stop RST segments from leaving a site to prevent them from reaching the attacker.

5.6.2 Reliability. In addition to being connectionless or connection-oriented, protocols are either *reliable* or *unreliable*. An unreliable protocol does not detect and correct errors. Some unreliable protocols do not even check for errors. Others check for errors but simply discard a message if they find that it contains an error.

TCP is a reliable protocol. It actually corrects errors. The TCP *checksum* field is calculated using values from other fields. The sender places the result of its calculation in the checksum field. The receiver redoes the calculation and compares it with the transmitted value. If the receiving transport layer process finds that a message is correct (the values are the same), it sends an acknowledgment message. However, if the receiver detects an error in the TCP segment it receives (the values are different), it discards the segment and does nothing else.

How does a receiver know that there is an error in the message? The sender computes a value based on the other bits in the TCP segment (not just the header). The receiver redoes the calculation. If the two values match, the receiver sends an acknowledgment. If they do not match, the receiver merely drops the segment and does not send an acknowledgment.

If the segment arrives correctly, the original sender receives an acknowledgment. However, if the segment never arrives or is discarded because of damage, no reply is sent. If the original sender does not receive an acknowledgment in a specified period of time, it will resend the original segment. It will even use the original sequence number.

5.6.3 Flag Fields. *Flag* field is a general name for a 1-bit field that is logical (true or false). To say that a flag field is *set* means that its value is 1. To say that a flag field is *not set* means that its value is 0.

The TCP header contains a number of flag fields. One of these is *SYN*. To request a connection opening, the sender sets the SYN bit. The other sends a *SYN/ACK* segment, in which both the SYN and ACK bits are set. Other commonly used flags are *FIN*, *RST*, *URG*, and *PSH*.

The URG flag indicates the presence of urgent data that should be handled before earlier data octets. The urgent pointer field indicates the location of the urgent data.

If an application message is large, TCP will divide the application message into multiple TCP segments and send the segments individually. To help the receiving TCP process, the sending transport process may set the PSH (*push*) bit in the application message's last segment. This tells the receiving transport process to push the data up to the application program immediately without buffering and delays.

5.6.4 Octets and Sequence Number. The *sequence number* field value allows the receiver to put arriving TCP segments in order even if the packets carrying them arrive out of order (including when a segment is retransmitted). Sequence numbers are also used in acknowledgments, albeit indirectly. In TCP transmission, every octet that is sent, from the very first, is counted. This octet counting is used to select each segment's sequence number.

- For the first segment, a random initial sequence number (*ISN*) is placed in the sequence number field.
- If the segment contains data, the number of the first octet contained in the data filed is used as the segment's sequence number.
- For a purely supervisory message that carries no data, such as an ACK, SYN, SYN/ACK, FIN, or RST segment, the sequence number is increased by 1 over the previous message.

One dangerous attack is TCP *session hijacking*, in which an attacker takes over the role of one side. This allows the hijacker to read messages and send false messages to the other side. To accomplish session hijacking, the attacker must be able to predict sequence numbers because if a segment arrives with an inappropriate sequence number, the receiver will reject it. TCP session hijacking is likely to be successful only if the initial sequence number is predictable. Few operating systems today pick initial sequence numbers in a predictable way, but predicable sequence numbers were common in earlier operating systems, some of which are still in use.

5.6.5 Acknowledgment Numbers. When a receiver sends an acknowledgment, it sets the ACK bit. It also puts a value in the *acknowledgment number* field to indicate which segment is being acknowledged. This process is needed because the sender sends many segments and because acknowledgments may be delayed.

You might think that the acknowledgment number would be the sequence number of the segment being acknowledged. Instead, it is the number of the last octet in the data field plus 1. In other words, the acknowledgment number gives the octet number of the first octet in the next segment to be sent. This seems a bit odd, but it makes certain calculations easier for the receiver.

5.6.6 Window Field. Flow control limits the rate at which a side sends TCP segments. The TCP *window* field allows one to limit how many more octets the other side may send before getting another acknowledgment. The process is somewhat complex and has no known security implications at the time of this writing (June 2013). In acknowledgments, the ACK bit is set, and both the acknowledgment and window size fields are filled in.

5.6.7 Options. Like the IPv4 header, the TCP header can have options. However, while IP options are rare and cause for suspicion, TCP uses options extensively. One common option, often sent with the initial SYN or SYN/ACK segment, is the *maximum segment size* (MSS) option. This gives the other side a limit on the maximum size of TCP segment data fields (not on segment sizes as a whole). The presence of TCP options, then, is not suspicious by itself.

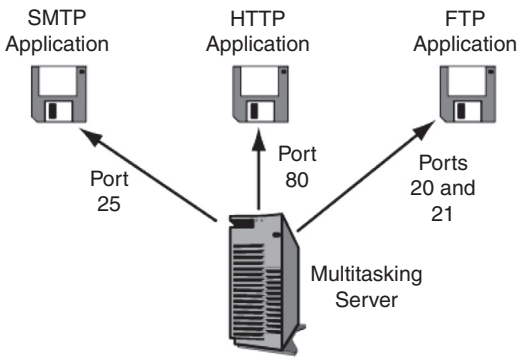


EXHIBIT 5.16 Multitasking Server Host and Port Numbers

5.6.8 Port Numbers. We have now looked at most fields in the TCP header. The first two fields warrant special mention.

5.6.8.1 Port Numbers on Servers. *Port number* fields mean different things for clients and servers. For a server, it represents a specific application running on that server, as Exhibit 5.16 shows. Servers are multitasking computers, which means that they can run multiple applications at the same time. Each application is specified by a different port number.

For instance, on a server, a Web server application program may run on TCP Port 80. Incoming TCP segments that have 80 as their destination port number are passed to the Web server application. Actually, TCP Port 80 is the well-known port for Web server programs, meaning that it is the usual port number for the application. Although Web servers can be given other TCP port numbers, this makes it impossible for users to establish connections unless they know or can guess the nonstandard TCP port number.

The TCP port range from 0 to 1023 is reserved for the well-known port numbers of major applications, such as HTTP and email. For instance, *Simple Mail Transfer Protocol* (SMTP) mail server programs usually are run on TCP Port 25, while *File Transfer Protocol* (FTP) requires two well-known port numbers—TCP Port 21 for supervisory control and TCP Port 20 for the actual transfer of files.

5.6.8.2 Port Numbers on Clients. Client hosts use TCP port numbers differently. Whenever a client connects to an application program on a server, it generates a random *ephemeral port number* that it uses only for that connection. On Windows machines, the ephemeral TCP port numbers range from 1024 to 4999.

The Microsoft port number range for ephemeral port numbers may differ from the official IETF range, with values of 5000–65534. The use of nonstandard ephemeral port numbers by Windows and some other operating systems causes problems for firewall filtering.

5.6.8.3 Sockets. Exhibit 5.17 shows that the goal of internetworking is to deliver application messages from one application on one machine to another application on another machine. On each machine, there is a TCP port number that specifies the application (or connection) and an IP address to specify a computer. A *socket* is a combination of an IP address and a TCP port number. It is written as the IP address,

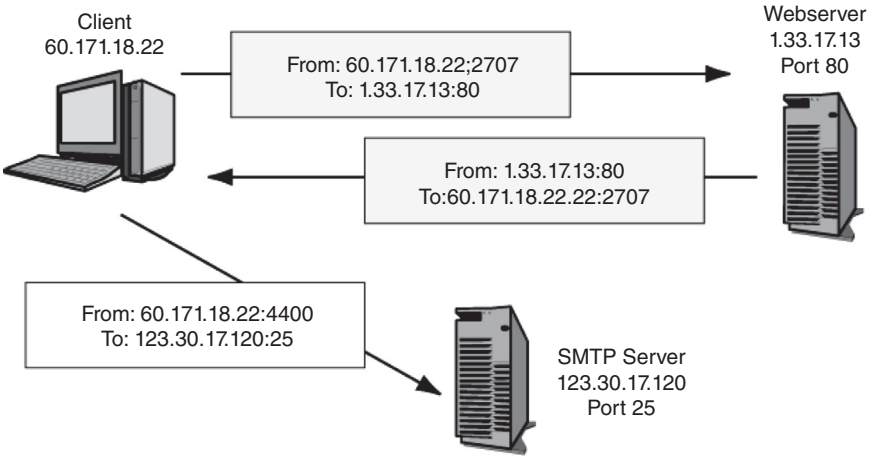


EXHIBIT 5.17 Sockets

a colon, and the TCP port number. A typical socket, then, would be something like *128.171.17.13:80*.

Attackers often do *socket spoofing*—both IP address spoofing and port spoofing. For instance, in TCP session hijacking, if the attacker wishes to take over the identity of a client, it must know both the client's IP address and ephemeral port number. Of course, these fields are transmitted in the clear (without encryption) in TCP, so an attacker with a sniffer that captures and reads traffic flowing between the client and server can easily obtain this information.

5.6.9 TCP Security. Like IP, TCP was created without security. However, although IPsec has made IP secure, the IETF has not created a comparable way to secure TCP. One reason for this is IPsec's ability to secure all transport layer traffic transparently, without modification to transport layer protocols. The IETF has made IPsec the centerpiece of its security protections and a single method to handle upper-layer security. Communicating partners that want TCP security should implement IPsec.

However, few TCP sessions are protected by IPsec. Consequently, some pairs of users employ an option in TCP, which adds an electronic signature to each TCP session. This signature proves the identity of the sender. This option, described in RFC 2385, requires the two parties to share a secret value. This option is awkward because it provides no way to share keys automatically, and it does not provide encryption or other protections. The option is used primarily in the Border Gateway Protocol (BGP). BGP is used to exchange routing information between administrative systems—say a corporate system and an internet service provider. BGP always uses one-to-one connections, the communicating parties usually know each other quite well, and the two parties have long-term relationships, which makes key exchange less burdensome and risky. Outside of BGP, however, the RFC 2385 electronic signature option does not appear to be used significantly. Even in BGP, it is widely seen as very weak security.

5.7 USER DATAGRAM PROTOCOL.

As noted earlier, TCP is a protocol that makes up for the limitations of IP. TCP adds error correction, the sequencing of IP packets, flow control, and other functionality that we have not discussed.

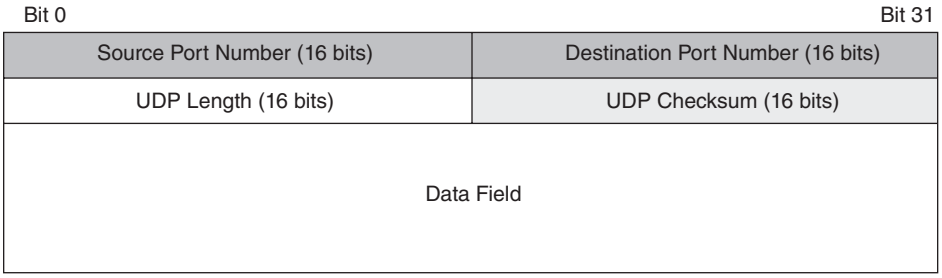


EXHIBIT 5.18 User Datagram Protocol (UDP)

Not all applications need the reliable service offered by TCP. For instance, in voice over IP (VOIP), there is no time to wait for the retransmission of lost or damaged packets carrying voice. In turn, the *Simple Network Management Protocol* (SNMP), which is used for network management communications, sends so many messages back and forth that the added traffic of connection-opening packets, acknowledgments, and other TCP supervisory segments could overload the network. Consequently, voice over IP, SNMP, and many other applications do not use TCP at the transport layer.

Instead, they use the *User Datagram Protocol* (UDP). This protocol is connectionless and unreliable. Each UDP message (called a UDP *datagram*) is sent on its own. There are no openings, closings, or acknowledgments.

As a consequence of the simplicity of UDP's operation, the UDP datagram's organization is also very simple, as Exhibit 5.18 illustrates. There are no sequence numbers, acknowledgment numbers, flag fields, or most of the other fields found in TCP.

There are source and destination port numbers, a UDP header length to allow variable-length UDP datagrams, and a UDP checksum. If the receiver detects an error using the checksum, it simply discards the message. There is no retransmission.

The fact that both TCP and UDP use port numbers means that whenever you refer to port numbers for well-known applications, you also need to refer to whether the port numbers are TCP or UDP port numbers. This is why the well-known port number for Web servers is TCP port 80.

TCP's sequence numbers make TCP session hijacking very difficult. The receiver will discard messages with the wrong sequence numbers even if the source and destination sockets are correct. UDP lacks this protection, making UDP a somewhat more dangerous protocol than TCP.

Like TCP, UDP has no inherent security. Companies that wish to secure their UDP communication must use IPsec.

5.8 TCP/IP SUPERVISORY STANDARDS.

So far, we have looked at standards that deliver a stream of packets across an internet and that perhaps check for errors and provide other assurances. However, the TCP/IP architecture also includes a number of supervisory protocols that keep the Internet functioning.

5.8.1 Internet Control Message Protocol (ICMP).

The first supervisory protocol on the Internet was the Internet Control Message Protocol (ICMP). As Exhibit 5.19 shows, ICMP messages are delivered in the data fields of IP packets.

The best-known pair of ICMP message types is the ICMP *echo* message and the *echo reply* message. Suppose that a host sends an ICMP echo message to an IP address. If a host is active at that address, it may send back an ICMP echo reply message. This

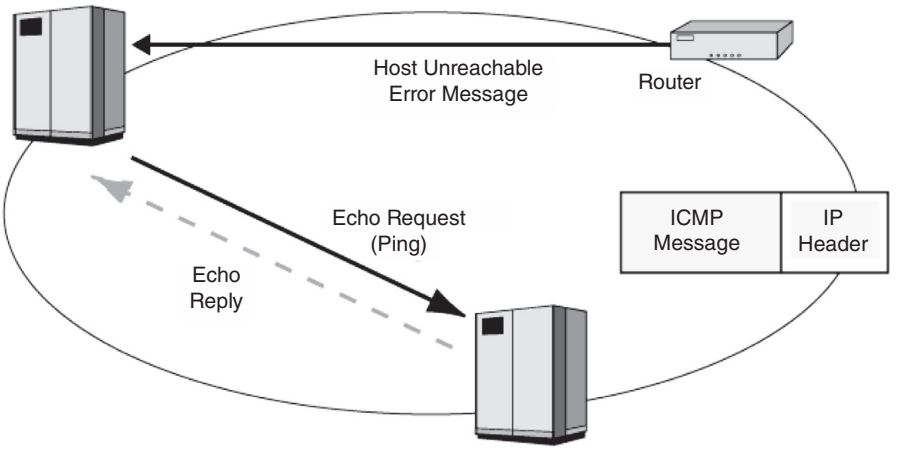


EXHIBIT 5.19 Internet Control Message Protocol (ICMP)

process is often called *pinging* because the most popular program for sending ICMP echo message is called *Ping*. The echo message is a very important tool for network management. If the network manager suspects a problem, he or she will ping a wide range of host addresses to see which of them are reachable. The pattern of responses can reveal where problems exist within a network.

Attackers also love to ping a wide range of host IP addresses. This can give them a list of hosts that are reachable for attacks. Another popular network management and attack tool is *traceroute* (or *tracert* on Windows PCs). *Traceroute* is similar to ping, but traceroute also lists the routers that lie between the sending host and the host that is the target of the traceroute command. This allows an attacker to map the network. Border firewalls often drop echo reply messages leaving the firm to the outside.

Many ICMP messages are error messages. For instance, if a router cannot deliver the packet, it may send back an ICMP error message to the source host. This error message will provide as much information as possible about the type of error that occurred.

If an attacker cannot ping destination hosts because a firewall stops them, attackers often send IP packets that are malformed and so will be rejected. The ICMP error message is delivered in an IP packet, and the source IP address in this packet will reveal the IP address of the sending router. By analyzing error messages, the attacker can learn how routers are organized in a network. This information can be very useful to attackers.

5.8.2 Domain Name System (DNS). To send a packet to another host, a source host must place the destination host's IP address in the destination address field of the packets. Often, however, the user merely types the host name of the destination host, for instance, cnn.com.

Unfortunately, host names are only nicknames. If the user types a host name, the computer must learn the corresponding IP address. As Exhibit 5.20 shows, the host wishing to send a packet to a target host sends a *Domain Name System* (DNS) request message to the DNS server. This message contains the host name of the target host. The DNS response message sends back the target host's IP address. To give an analogy, if you know someone's name, you must look up their telephone number in a telephone directory if you want to call them. In DNS, the human name corresponds to the

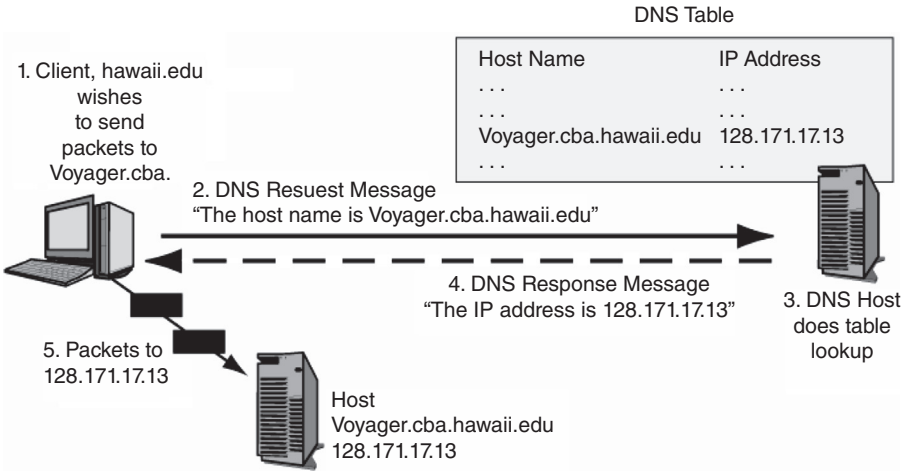


EXHIBIT 5.20 Domain Name System (DNS) Server

host name, the telephone number corresponds to the IP address, and the DNS server corresponds to the telephone directory.

DNS is critical to the Internet's operation. Unfortunately, DNS is vulnerable to several attacks. For example, in *DNS cache poisoning*, an attacker replaces the IP address of a host name with another IP address. After cache poisoning, a legitimate user who contacts a DNS server to look up the host name will be given the false IP address, sending the user to the attacker's chosen site. Denial-of-service attacks are also too easy to accomplish. RFC 3833 lists a number of DNS security issues.[2]

Several attempts to strengthen DNS security have been developed, under the general banner of Domain Name System Security Extensions (*DNSSEC*), especially RFC 2535.[3] However, both the original DNSSEC specifications and the newer DNSSEC *bis* specifications (RFCs 4033-4035[4]) have proven to be insufficient. Developing a security standard that is sufficiently backwardly compatible for Internet-scale implementation has proven to be extremely difficult.

If the DNS server does not know the host name, it contacts another DNS server. The DNS system contains many DNS servers organized in a hierarchy. At the top of the hierarchy are 13 *DNS root servers*. Below these are DNS servers for *top-level domains*, such as *.com, .edu, .ie, .uk, .nl,* and *.ca.* Each top-level domain has two or more top-level DNS servers for their domain. Second-level domain names are given to organizations (e.g., *Hawaii.edu* and *Microsoft.com*). Organizations are required to maintain DNS servers for computers within their domain.

If attackers could bring down the 13 root servers, they could paralyze the Internet. Widespread paralysis would not occur immediately, but in a few days, the Internet would begin experiencing serious outages.

5.8.3 Dynamic Host Configuration Protocol (DHCP). Server hosts are given static (permanent) IP addresses. Client PCs, however, are given dynamic (temporary) IP addresses whenever they use the Internet. The *Dynamic Host Configuration Protocol* (DHCP) standard that we saw earlier in the chapter makes this possible. A DHCP server has a database of available IP addresses. When a client requests an IP address, the DHCP server picks one from the database and sends it to the client. The next time the client uses the Internet, the DHCP server may give it a different IP address.

The fact that clients may receive different IP addresses each time they get on the Internet causes problems for *peer-to-peer* (P2P) applications. A presence server or some other mechanism must be used to find the other party's IP address. A lack of accepted standards for presence (including presence security) is a serious issue now that P2P applications are widespread. In fact, most security considerations in P2P presence servers have been used in P2P piracy applications, with an eye toward avoiding discovery by legitimate authorities.

5.8.4 Dynamic Routing Protocols. How do routers on the Internet learn what to do with packets addressed to various IP addresses? They frequently talk to one another, exchanging information about the organization of the Internet. These exchanges must occur frequently because the structure of the Internet changes frequently as routers are added or dropped. Protocols for exchanging organization information are called *dynamic routing protocols*. There are many dynamic routing protocols, including the *Routing Information Protocol* (RIP), *Open Shortest Path First* (OSPF), the *Border Gateway Protocol (*BGP), and Cisco Systems' proprietary *Enhanced Interior Gateway Routing Protocol* (EIGRP). Each is used under different circumstances. These protocols have widely different security features, and different versions of each protocol have different levels of functionality.

An attacker who can impersonate a router can send false dynamic routing protocol messages to other routers. These false messages could cause the routers to fail to deliver their packets. The attacker could even cause packets to pass through the attacker's computer (called a *man-in-the-middle attack* or MIMA) in order to read their contents.

The protocols just listed have widely different security features, and different versions of each protocol have different levels of security functionality.

5.8.5 Simple Network Management Protocol (SNMP). Networks often have many elements—routers, switches, and host computers. Managing dozens, hundreds, or thousands of devices can be nearly impossible. To make management easier, the IETF developed the *Simple Network Management Protocol* (SNMP). As Exhibit 5.21 shows, the manager program can send SNMP messages to managed devices to determine their conditions. The manager program can even send configuration messages that can change the ways in which remote devices operate. This allows the manager to fix many problems remotely.

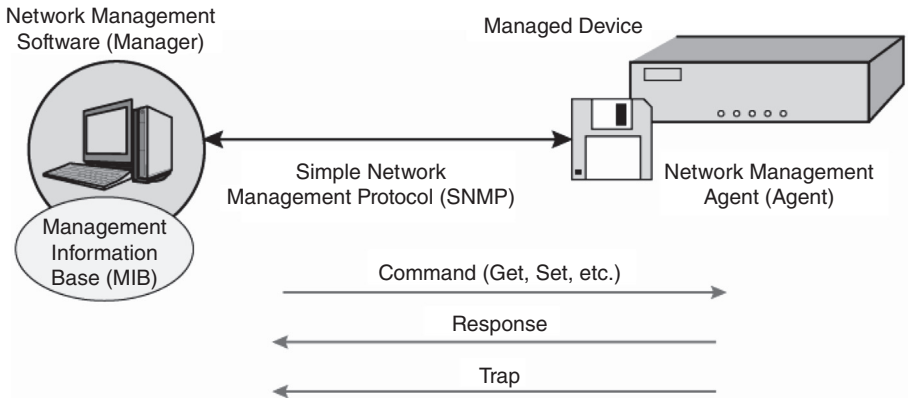


EXHIBIT 5.21 Simple Network Management Protocol (SNMP)

Many firms disable remote configuration because of the damage that attackers could do with it. They could simply turn off all ports on switches and routers, or they could do more subtle damage.

5.9 APPLICATION STANDARDS. Most applications have their own application-layer standards. In fact, given the large number of applications in the world, there are literally hundreds of application-layer standards.

As corporations get better at defending against attacks at lower layers, attackers have begun to focus their attention on application vulnerabilities. If an attacker can take over an application running with high privileges, he or she obtains these privileges. Many applications run at the highest privileges, and attackers that compromise them own the box.

5.9.1 HTTP and HTML. Many applications have two types of standards. The *transport standard* transfers application-layer messages between applications on different machines; for the World Wide Web, this is the *Hypertext Transfer Protocol* (HTTP). The other is a standard for *document structure*. The main document-structure standard for the WWW is the *Hypertext Markup Language* (HTML).

Netscape, which created the first widely used browser, also created a security standard to protect HTTP communication. This was *Secure Sockets Layer* (SSL). Later, the Internet Engineering Task Force took over SSL and changed the name of the standard to *Transport Layer Security* (TLS).

5.9.2 E-Mail. Popular transfer standards for email are the *Simple Mail Transfer Protocol* (SMTP), *Post Office Protocol* (POP), and *Internet Message Access Protocol* (IMAP) for downloading email to a client from a mailbox on a server. Popular document-body standards include RFC 2822 (for all-text messages), HTML, and Multipurpose Internet Mail Extensions (MIME). S/MIME (Secure MIME) adds public-key encryption (see Chapter 7) to MIME and is defined in RFCs 2634, 3850, and 3851.

An obvious security issue in email is content filtering. Viruses, spam, phishing messages, and other undesirable content should be filtered out before they reach users and can do damage. (For more information on spam and other low-technology attacks, see Chapter 20 in this *Handbook*; for malware and spam countermeasures see Chapters 26, 27, 31, and 41.)

Another security issue in email is securing messages flowing from the sending client to the sender's mail server, to the receiver's mail server, and to the receiving client. Fortunately, there are security standards for part or all of the message flows, including SSL/TLS and S/MIME among others. Unfortunately, the IETF has been unable to agree on a security standard.

When Web mail, which uses HTTP and HTML for email communication, is used, then SSL/TLS can work between the sender and the sender's mail server and between the receiver's mail server and the receiver. Transmission between the email servers is another issue. Of course, senders can send encrypted message bodies directly to receivers. However, this prevents filtering at firewalls. Users should be particularly careful about using Web mail via wireless connections. (See Chapters 32 and 33.)

5.9.3 Telnet, FTP, and SSH. The two earliest applications on the Internet were the *File Transfer Protocol* (FTP) and *Telnet*. FTP provides bulk file transfers between hosts. Telnet allows a user to launch a command shell (user interface) on another computer. Neither of these standards has any security. Of particular concern is that

both send passwords in the clear (without encryption) during login. The newer *Secure SHell* (SSH) standard can be used in place of both FTP and Telnet while providing high security by encrypting all transferred traffic between the hosts.

5.9.4 Other Application Standards. There are many other applications and therefore application standards. These include *Voice over IP* (VoIP; see Chapter 34 in this *Handbook*), *peer-to-peer applications* (P2P; see Chapter 35), and *service-oriented architecture* (SOA) and Web service applications (see Chapters 21, 30, and 31), among many others. Most applications have serious security issues. Application security has become perhaps the most complex aspect of network security (see Chapters 38, 39, and 40).

5.10 CONCLUDING REMARKS. It is impossible to understand information security without a strong knowledge of networking. This chapter is designed to give you a working overview of networking. It is likely to be sufficient if you run into basic networking questions while reading other chapters in this *Handbook*. However, to work in security, you will need a much stronger knowledge of networking. The books and other resources cited in Section 5.11 are a good start in that direction.

5.11 FURTHER READING

Comer, D. E. *Internetworking with TCP/IP Vol. 1: Principles, Protocols, and Architecture*, 6th ed. Addison-Wesley, 2013.

Ferrero, A. *The Eternal Ethernet*, 2nd ed. Boston: Addison-Wesley, 1999.

FitzGerald, J. *Business Data Communications and Networking*, 11th ed. Wiley, 2011.

Freedman, A. *Computer Desktop Encyclopedia*. Point Pleasant, PA: Computer Language Company, 2013. Available online from www.computerlanguage.com

Gibson, D. *Microsoft Windows Networking Essentials*. Sybex, 2011

Hallberg, B. *Networking, A Beginner's Guide*, 5th ed. McGraw-Hill Osborne Media, 2009.

Hummel, S. L. *Network Design Fundamentals*. CreateSpace Independent Publishing Platform, 2013.

Kurose, J. F., and K. W. Ross. *Computer Networking: A Top-Down Approach*, 6th ed. Pearson, 2012.

Panko, R. R. *Business Data Networks and Telecommunications*, 8th ed. Prentice-Hall, 2010.

Panko, R. R. *Corporate Computer and Network Security*, 2nd ed. Prentice-Hall, 2009.

Panko, R. R., and J. Panko. *Business Data Networks and Security*, 9th ed. Prentice-Hall, 2012.

Palmer, M. *Hands-On Networking Fundamentals*, 2nd ed. Cengage Learning, 2012.

Roberts, R. M. *Networking Fundamentals*, 2nd ed. Goodheart-Willcox, 2011.

Spurgeon, C. *Ethernet: The Definitive Guide*. O'Reilly Media, 2000.

Stevens, W. R. *TCP/IP Illustrated, Volume 1: The Protocols*, 2nd ed. Addison-Wesley Professional, 2011.

5.12 NOTES

1. P. Prabakaran, "Tutorial on Spread Spectrum Technology," *EE Times | Design,* May 6, 2003, www.eetimes.com/design/communications-design/4008962/Tutorial-on-Spread-Spectrum-Technology
2. www.faqs.org/rfcs/rfc3833.html
3. www.faqs.org/rfcs/rfc2535.html
4. www.dnssec.net/

CHAPTER 6

LOCAL AREA NETWORK TOPOLOGIES, PROTOCOLS, AND DESIGN

Gary C. Kessler

6.1 OVERVIEW 6·2
- 6.1.1 LAN Characteristics 6·2
- 6.1.2 LAN Components 6·2
- 6.1.3 LAN Technology Parameters 6·3
- 6.1.4 Summary 6·3

6.2 LAN TOPOLOGY 6·3
- 6.2.1 Network Control 6·4
- 6.2.2 Star Topology 6·4
- 6.2.3 Ring Topology 6·4
- 6.2.4 Bus Topology 6·6
- 6.2.5 Physical versus Logical Topology 6·6

6.3 MEDIA 6·8
- 6.3.1 Coaxial Cable 6·8
- 6.3.2 Twisted Pair 6·9
- 6.3.3 Optical Fiber 6·10
- 6.3.4 Wireless Media 6·11
- 6.3.5 Summary 6·13

6.4 MEDIA ACCESS CONTROL 6·13
- 6.4.1 Contention 6·13
- 6.4.2 Distributed Polling 6·14

6.5 LAN PROTOCOLS AND STANDARDS 6·15
- 6.5.1 OSI Model versus LAN Model Architectures 6·15
- 6.5.2 IEEE 802 Standards 6·17
- 6.5.3 IEEE 802.3 CSMA/CD Standard 6·19
- 6.5.4 Ethernet II 6·21
- 6.5.5 IEEE 802.5 Token-Ring Standard 6·22
- 6.5.6 IEEE 802.2 LLC Standard 6·23
- 6.5.7 Summary 6·24

6.6 INTERCONNECTION DEVICES 6·24
- 6.6.1 Hubs 6·25
- 6.6.2 Switches 6·25
- 6.6.3 Bridges 6·25
- 6.6.4 Routers 6·26
- 6.6.5 Summary 6·27

6.7 NETWORK OPERATING SYSTEMS 6·27

6.8 SUMMARY 6·28

6.9 FURTHER READING 6·30

6.10 NOTES 6·30

This chapter provides a broad overview of local area network (LAN) concepts, basic terms, standards, and technologies. These topics are important to give the information security professional a better understanding of the terms that might be used to describe a particular network implementation and its products. The chapter also is written with an eye to what information security professionals need to know; for a more complete

overview of the topic, the reader is referred to general LAN texts, such as those listed in Section 6.10.

6.1 OVERVIEW. There are a number of ways to describe a LAN, and each will provide a glimpse as to implementation and product differences as well as points of security exposures. This section introduces various terms and perspectives as a basis for the discussion in the following sections.[1]

6.1.1 LAN Characteristics. One way of describing LANs is to describe the characteristics that distinguish a *local* network from other types of networks. The most common characteristics are:

- Small geographic scope (the two most distant stations may be up to 5 kilometers [km] or so apart)
- Fast speed (data rates well in excess of 1 million [mega] bits per second [Mbps] and up to 1 billion [giga] bits per second [Gbps])
- Special media (common use of coaxial cable and optical fiber, as well as twisted pair)
- Private ownership

This type of network, then, has a very different look and feel than the Internet or some other public or private wide-area networks (WANs). More people have access to the LAN infrastructure than to the infrastructure of just about any WAN. LAN users can easily "spy" on each other by sniffing packets, something that is generally very difficult on the Internet. A single user can bring the LAN to a standstill.

The corporate LAN is generally the users' primary access to the Internet. The users on the LAN are behind the corporate firewall and router; some studies suggest that they are responsible for 80 percent of security incidents.

Other LANs include hotel networks, which are often used by criminals as handy sources of confidential information available on unprotected systems temporarily hooked up to these Internet-access services.

Often the success of these attacks is due to users' lack of education and awareness, such as choosing poor passwords, not maintaining up-to-date virus signature files, or computers attached to the LAN without firewalls. Sometimes the attacks are more sophisticated, such as using a packet sniffer to learn another user's password or taking steps to degrade network performance.

6.1.2 LAN Components. In general, there are four basic components required to build a LAN, providing their own vulnerabilities and exposures from a security perspective:

1. **Computers**. These are the basic devices that are connected on the network. Read "computer" very broadly; the term can include personal computers (PCs), minicomputers, mainframes, file servers, printers, plotters, mobile devices (e.g., smartphones and tablets), communications servers, and network interconnection devices. It can also include protocol analyzers.
2. **Media**. These are the physical means by which the computers are interconnected. LAN media include unshielded twisted pair (UTP), coaxial cable (coax), optical fiber, and wireless (radio) devices. The wireless media have connection points

throughout an area where devices can attach to the network, and every place is a potential connection point in a wireless environment.

3. **Network interface card (NIC)**. This is the physical attachment from the computer to the LAN medium. Older NICs are internal cards; the only item that is actually seen is the physical attachment to the LAN, often an RJ-45 jack. An increasing number of adapters use the universal serial bus (USB) slots on modern computers. Although NICs range widely in price depending on their capabilities, intended use, and vendor, an internal 1 Gbps Ethernet NIC for a desktop personal computer (PC) could be purchased for less than $7, a USB Ethernet adapter for about $14, and a wireless USB adapter for less than $13 at the time of writing in January 2013.
4. **Software**. The three components above provide physical connectivity. Software—often called a *network operating system* (NOS)—is necessary for the devices to actually take advantage of the resource sharing that the LAN can provide. The NOS can support many types of services such as file sharing, print sharing, client/server operation, communications services, and more.

While the LAN needs to be examined in a holistic fashion, each of these components at each attached node also may require examination.

6.1.3 LAN Technology Parameters. One final way of discussing the specific operation of the LAN is to describe the technology:

- **Physical topology**. The physical layout of the medium.
- **Logical topology**. The logical relationship of the LAN nodes to each other.
- **Media Access Control (MAC) Standard**. The specification describing the rules that each node follows to determine when it is its turn to transmit on the medium.
- **Use of the Logical Link Control (LLC) protocol**. Defines the frame format employed above the MAC layer, and additional services.
- **Use of higher-layer protocols**. Defines the node-to-node communicating protocols and additional higher-layer applications.

6.1.4 Summary. It does not matter how a LAN is classified or described. It is essential, however, that the LAN be understood from a variety of perspectives to be able to apply a network-security examination.

6.2 LAN TOPOLOGY. WANs typically use some sort of switched technology, such as traditional circuit switching, packet switching (e.g., X.25), or fast packet switching (e.g., frame relay or asynchronous transfer mode [ATM]). Indeed, point-to-point lines typically connect the network switches so that there is a single data transmission on the line at one time.

Historically, LANs have been broadcast networks, meaning that every LAN station hears every transmission on the medium. LAN topologies, then, have to support the broadcast nature of the network and provide full connectivity between all stations.

The *topology* of a network is used to describe two issues. The *physical topology* describes how the LAN stations are physically connected so that they can communicate with each other. The *logical topology* describes how the broadcast nature of the LAN is actually affected, and, therefore, how stations participate in the process of obtaining

permission to transmit on the medium. There are three common topologies found in LANs: star, ring, and bus.

6.2.1 Network Control. Since LANs are broadcast networks, it is imperative that only a single node be allowed to transmit at any one time. All LANs use a *distributed access-control* scheme, meaning that all nodes follow the same rules to access the network medium and no one LAN node controls the other nodes' access. In this way, LAN nodes can come online and offline without bringing the network down.

This description is *not* meant to suggest that there are no critical elements in a LAN. Indeed, if a central hub, switch, or transmitter fails, the LAN will crash. Distributed control *does* suggest, however, that all nodes (user stations) follow the same access rules, and failure of a single *node* will not bring the LAN down. The access-control scheme is defined by the MAC protocol.

6.2.2 Star Topology. In a *star topology* (see Exhibit 6.1), all devices on the LAN are interconnected through some central device. Since LANs use distributed access-control schemes, all communication is from one node to another, and the central device merely provides a pathway between pairs of devices.

Physical star topologies have a tremendous advantage over other topologies in that they greatly ease network administration, maintenance, reconfiguration, and error recovery. Disadvantages include the single point of failure.

6.2.3 Ring Topology. In a *ring topology*, the nodes are connected by a set of point-to-point links that are organized in a circle (see Exhibit 6.2). Stations connect to

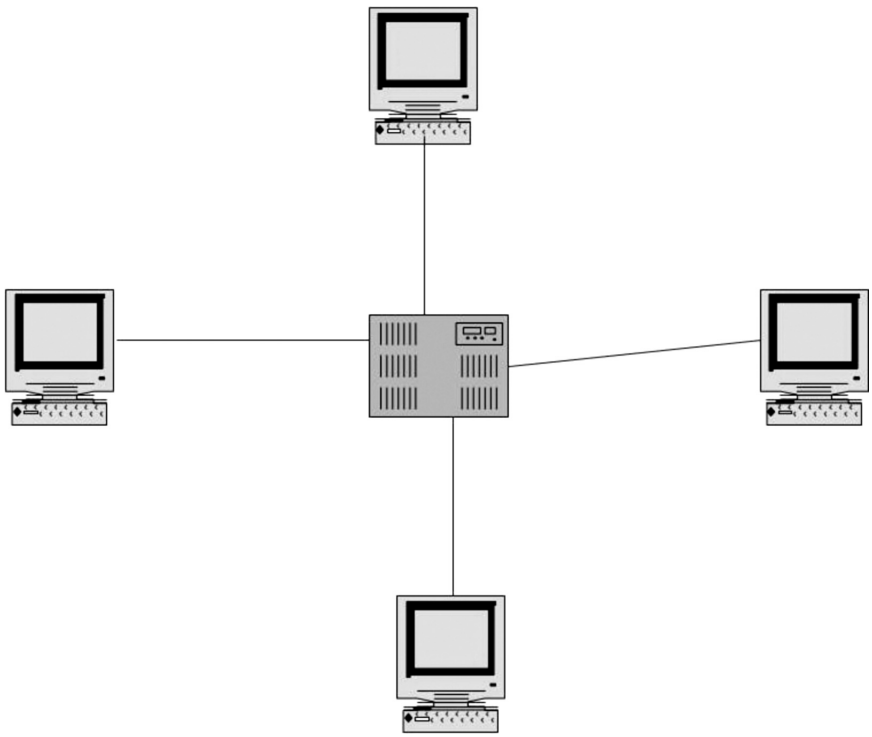


EXHIBIT 6.1 Star Topology

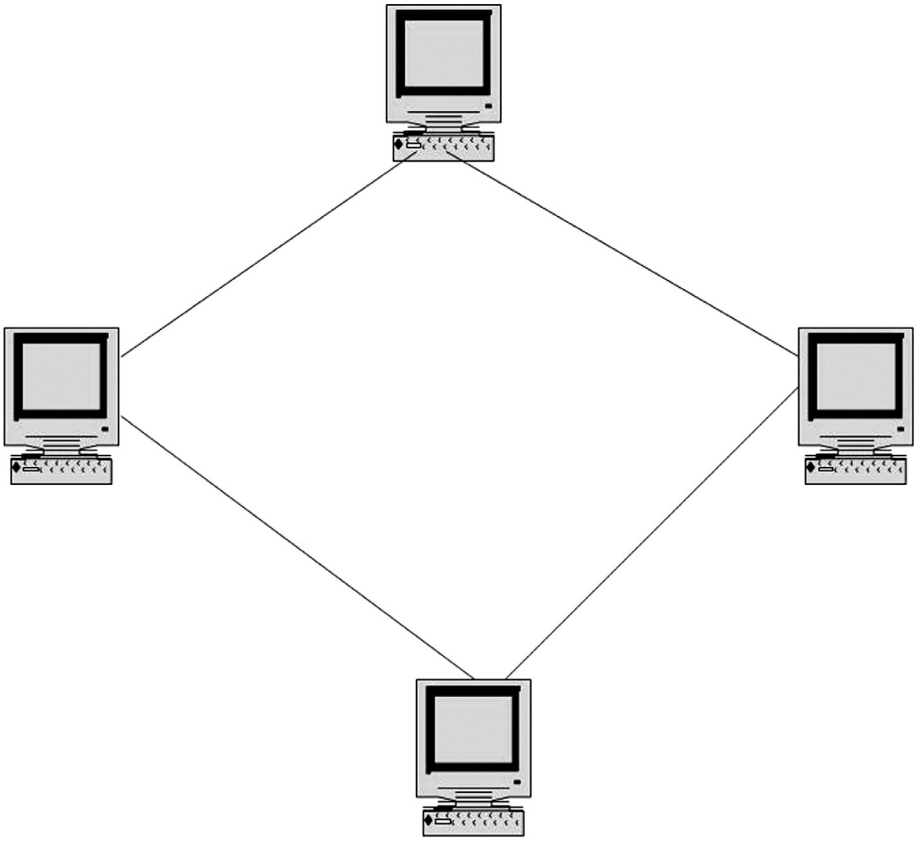


EXHIBIT 6.2 Ring Topology

the medium using *active taps* that are actually bit repeaters; a bit is read from the input line, held for a single bit time, then transmitted out to the output line.

A station transmits a message on the network by sending out a bit stream on its outgoing link; thus, rings are unidirectional in nature. Since all of the other stations see the bits one at a time, the intended receiver has no prior warning about an incoming message. For this reason, the transmitter is responsible for removing the message from the ring when the bits come back around. The MAC scheme ensures that multiple stations do not transmit at the same time.

In addition, a ring is a *serial broadcast* network. Because a station sends a message one bit at a time, every other station will see the message as it passes through but each will be receiving a different part of the message at any point in time.

Rings are a common physical LAN topology. However, unlike stars, they have multiple points of failure: if one link or one active tap fails, the integrity of the ring is destroyed. If the probability of failure of a single element is p and there are n elements, then the probability of failure $P\{F\}$ of the LAN is

$$P\{F\} = 1 - (1 - p)^n$$

As the number of elements rises, the probability of network failure rises exponentially. This problem is of such a critical nature that nearly all ring products use a star-wiring scheme or have some sort of redundancy built in for just this eventuality.

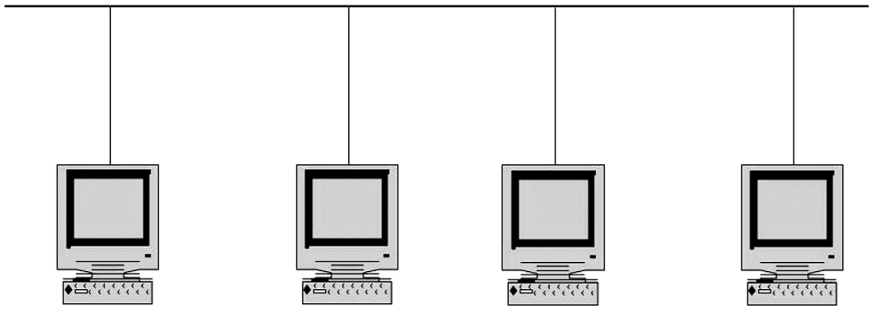


EXHIBIT 6.3 Bus Topology

6.2.4 Bus Topology. In a *bus topology* (see Exhibit 6.3), all devices are connected to a single electrically continuous medium; for this reason, this topology is also called a common cable or shared medium network. Nodes attach to the medium using a *passive tap*, one that monitors the bit flow without altering it. This is similar to the operation of a voltmeter; it measures the voltage on a power line without changing the available voltage.

Bus networks are analogous to the way appliances are connected to an alternating current (AC) power line. All of the devices draw power from the same source, even if they are on different physical segments of the power distribution network within the building. In addition, the operation of the devices is independent of each other; if the coffeepot breaks, the toaster will still work.

A bus is a *simultaneous broadcast* network, meaning that all stations receive a transmitted message at essentially the same time (ignoring propagation delay through the medium). Most home and business LANs employ a *baseband bus* where direct current (DC) signals are applied directly to the bus by the transmitter without any modification. In addition, transmissions on a baseband bus are broadcast bidirectionally and cannot be altered by the receivers. Bus LAN technologies are employed on cable television systems. For example, they employ a *broadband bus* where the signals are modulated (i.e., frequency shifted) to certain frequencies for transmission in one direction or another.

Buses are the oldest LAN topology and are generally limited in the type of medium that they can use. They do not usually suffer from single-point-of-failure problems.

6.2.5 Physical versus Logical Topology. A distinction was made above between the *physical* and *logical* topology of a LAN. *Physical* topology describes how the stations are physically positioned and attached to each other whereas the *logical* topology describes how the signals propagate and the logical operation of the network.

In all of today's commonly used LANs, the logical topology differs from the physical topology. The most common LAN configuration today is a *star-wired bus* (see Exhibit 6.4). This type of network has a star topology where all stations are physically attached with point-to-point links to a central device. This central device contains a bus that interconnects all of the I/O ports in such a way that when one station transmits a message, all stations will receive it. Since this acts exactly like a simultaneous broadcast, or bus, network, we categorize this configuration as a physical star, logical bus.

Another common configuration is a *star-wired ring* (see Exhibit 6.5). In this configuration, the bits will travel in logical order from station A to B, C, A, and so forth,

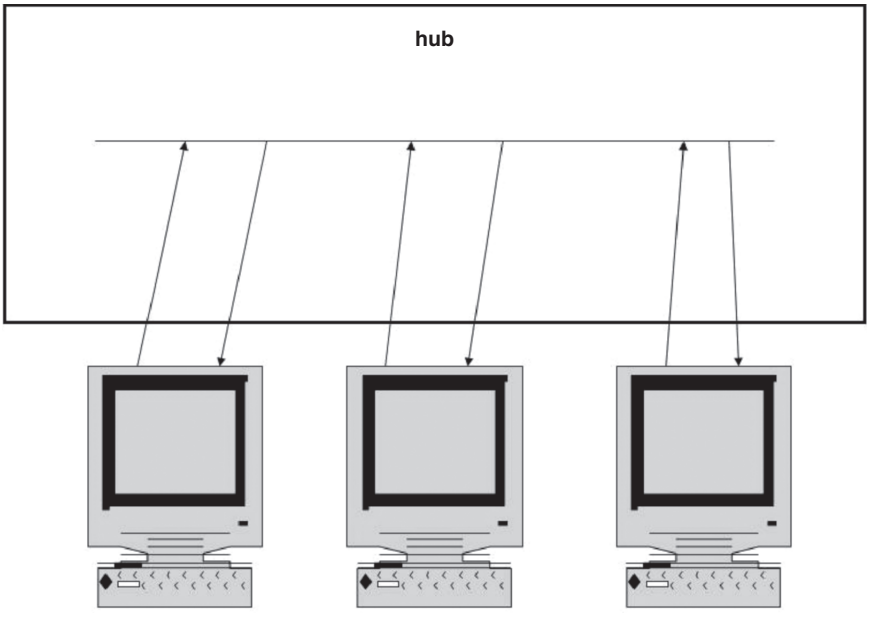


EXHIBIT 6.4 Star-Wired Bus

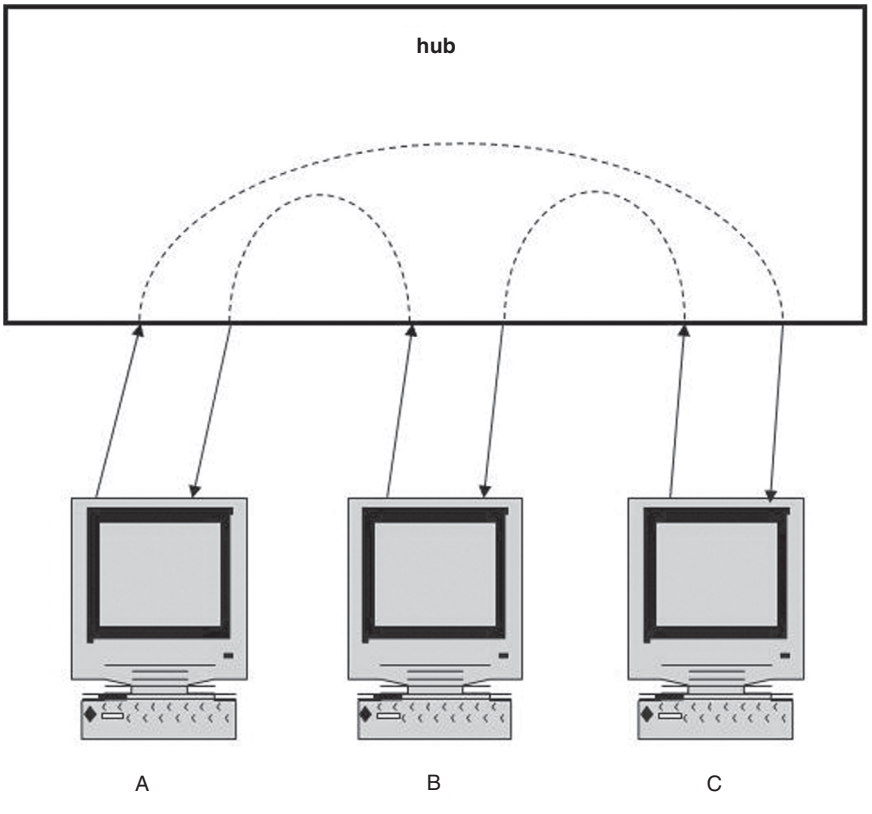


EXHIBIT 6.5 Star-Wired Ring

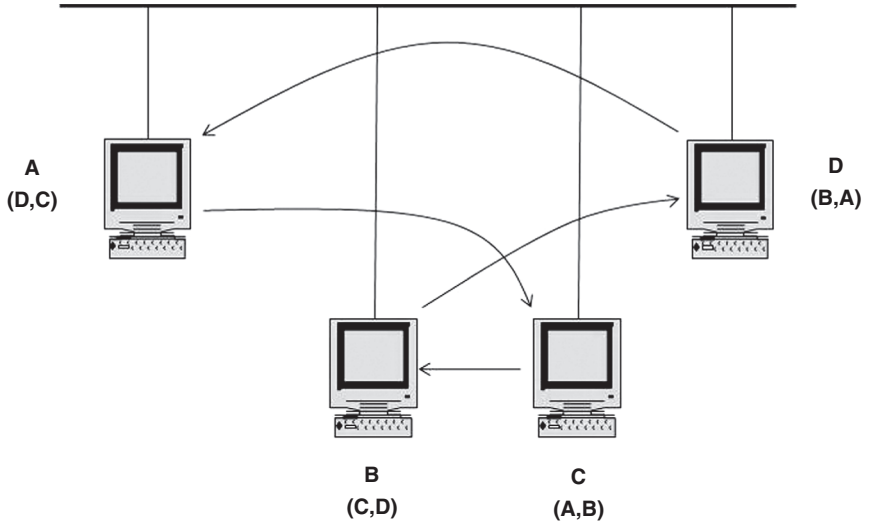


EXHIBIT 6.6 Bus-Wired Ring (The station identifier is shown above the station ID of the predecessor and successor stations in the logical ring.)

which matches the serial broadcast operation of a ring. We call this a physical star, logical ring.

Although uncommon today, another hybrid technology is the *bus-wired ring* (see Exhibit 6.6). In this configuration, nodes are passively attached to a single cable, forming a physical bus. Each station maintains a table specifying the address of predecessor and successor stations, thus forming a logical ring.

6.3 MEDIA. The next paragraphs discuss the three primary types of LAN media currently in use. Due to their relatively high speed, small geographic size, and protected environments, a number of media types can be employed with LANs.

6.3.1 Coaxial Cable. Coaxial cable (coax) is the original LAN medium. It gets its name from the physical composition of the cable itself (see Exhibit 6.7). At the center of the cable is a conductor, usually made of copper, which is surrounded by an insulator that, in turn, is surrounded by another conductor that acts as an electrical shield. Since the shield completely surrounds the central conductor and the two have a common axis, the shield prevents external electrical noise from affecting signals on

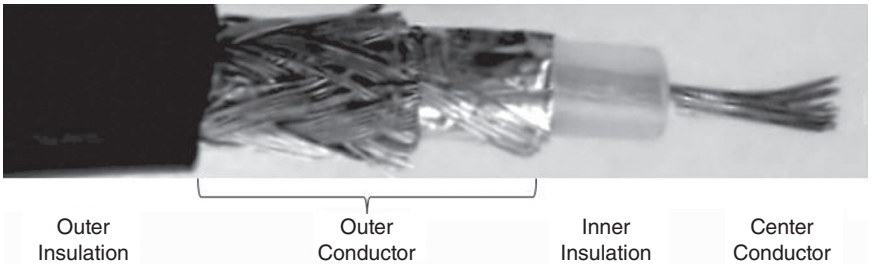


EXHIBIT 6.7 Coaxial Cable



EXHIBIT 6.8 Unshielded Twisted Pair

the conductor and prevents signals on the conductor from generating noise that affects other cables.

Coaxial cables vary in size from $\frac{1}{4}$–1 inch (6.35–25.4 mm), depending on the thickness of the conductor, shield, and insulation. Applications for coax range from cable television to LANs. Speeds in excess of several hundred Mbps at distances of several hundred to several thousand meters can be achieved. Coaxial cable also has a high immunity from electromagnetic and radio frequency interference. However, it is easy to *tap*.

Coaxial cable is only seen in physical bus LANs such as Ethernet. The original Ethernet specification, in fact, called for a thin coaxial cable; a later version that employed thin (CATV) coax was dubbed CheaperNet. Coax is not typically found in star or ring networks.

Coaxial cable is easy to wiretap.[2]

6.3.2 Twisted Pair. The medium enjoying the largest popularity for LAN applications today is twisted pair. Twisted pair cable consists of two insulated copper conductors that are twisted around each other (see Exhibit 6.8). This is typically 22- to 26-gauge (i.e., 0.025"/0.644 mm to 0.016"/0.405 mm) wire, the same as is used for telephone wiring. Twisting the conductors around each other minimizes the effect of external electrical radiation on the signal carried on the wire; if external voltage is applied to one wire of the pair, it will be applied equally to the other wire. The twisting, then, effectively eliminates the effect of the external noise. As the number of twists per inch increases, the noise reduction characteristics improve; unfortunately, so does the overall amount of cable and the cost. Most twisted pair for telephony applications has 10 to 15 twists per foot.

The type of twisted pair cable shown in Exhibit 6.8 is called *unshielded twisted pair* (UTP) because the wire pair itself is not shielded. The data-carrying capacity of UTP is generally indicated by its *category*:

Level 1 (sometimes called Category 1, or Cat 1) cable is older 0.4 MHz cable used for some telephone and modem applications, but is generally unsuited for data applications.

Level 2 (sometimes called Category 2, or Cat 2) cable is 4 MHz cable used for legacy data terminal systems, such as IBM 3270 BISYNC.

Category 3 (Cat 3) cable has a maximum bandwidth of 16 MHz and is rated for 10 Mbps over a wire segment of 100 m (although speeds of 100 Mbps can often be achieved). Cat 3 cable is rated up to 16 Mbps and is primarily used today for telephones.

Category 4 (Cat 4) cable, with a maximum bandwidth of 20 MHz, was used for IBM's 16 Mbps Token Ring networks. It is not commonly seen today.

Category 5 (Cat 5) cable has a maximum bandwidth of 100 MHz and is rated for voice or data at speeds up to 100 Mbps over a wire segment of 100 meters. Cat 5e is rated for full-duplex and 1 gigabit (Gbps) Ethernet. These are the most common LAN cables in use today.

Category 6 (Cat 6) cable is rated to 250 MHz over a wire segment between 15 and 100 meters in length. Cat 6 is intended for use for very-high-speed broadband applications at data rates up to 10 Gbps. Cat 6a cable is a variant of Cat 6, rated to 500 MHz and 10 Gbps.

Category 7 (Cat 7) cable is rated up to 600 MHz and uses four pair. Each pair of wires in the cable sheath, and the sheath itself, are shielded to prevent electromagnetic interference at data rates up to 10 Gbps. Cat 7a is rated up to 1,000 MHz and 10 Gbps data rates.

UTP is commonly found in physical star-wired bus and ring LANs; it is never used in a physical bus and rarely in a physical ring. As with coaxial cable, it is easy to wiretap. In addition, many LANs using UTP also make connections through *patch panels*, which are frequently unprotected because the technicians installing the connections are unaware of the security issues of providing centralized access to dozens or hundreds of connections.

Another twisted pair variant is *shielded twisted pair* (STP), where each cable pair is surrounded by a metallic shield that provides the same function as the outer conductor in coaxial cable. STP was only used in the IBM Token Ring, a star-wired ring.

6.3.3 Optical Fiber. Optical fiber is a thin flexible medium that acts as a waveguide for signals in the 10^{14}- to 10^{15}-Hz range, which includes the visible light spectrum and part of the infrared spectrum. Optical fiber is a great medium for digital communications; it is essentially immune to any type of radio or magnetic interference and very difficult (using highly specialized equipment) to tap surreptitiously. Theoretically able to achieve data rates on the order of trillions of bits per second, optical fiber has been shown to reach data rates of 100 Gbps over a 4,350-mile (7,000-km) fiber; the practical limit is usually due to the electronics performing optical-electrical conversion.

In WAN applications, this speed limit is exceeded in one of two ways.

1. An optical switch can terminate optical fiber without any electrical-optical conversion.
2. Dense wave division multiplexing (DWDM) allows 100 or more 10 Gbps bit streams to be carried on a single-fiber strand simultaneously. These technologies may well eventually find their way to the LAN.

The electronics are a critical part of any optical fiber system. The incoming electrical signal to be transmitted on the fiber is converted to an optical signal by the transmitter. Common optical sources are a *light-emitting diode* (LED) or *injection laser diode* (ILD). LEDs are less expensive than ILDs but are limited to lower speeds. The optical signal is received by a device called a *photodiode*, which essentially counts photons and converts the count to an electrical signal. Common photodiodes include the *positive-intrinsic-negative* (PIN) *photodiode* and *avalanche photodiode* (APD). The PIN is less expensive than the APD but is limited to lower speeds.

The physical and transmission characteristics of optical fiber are shown in Exhibit 6.9. At the center of an optical fiber cable is the *core*, a thin, flexible medium capable of carrying a light signal. The core is typically between 2 and 125 micrometers

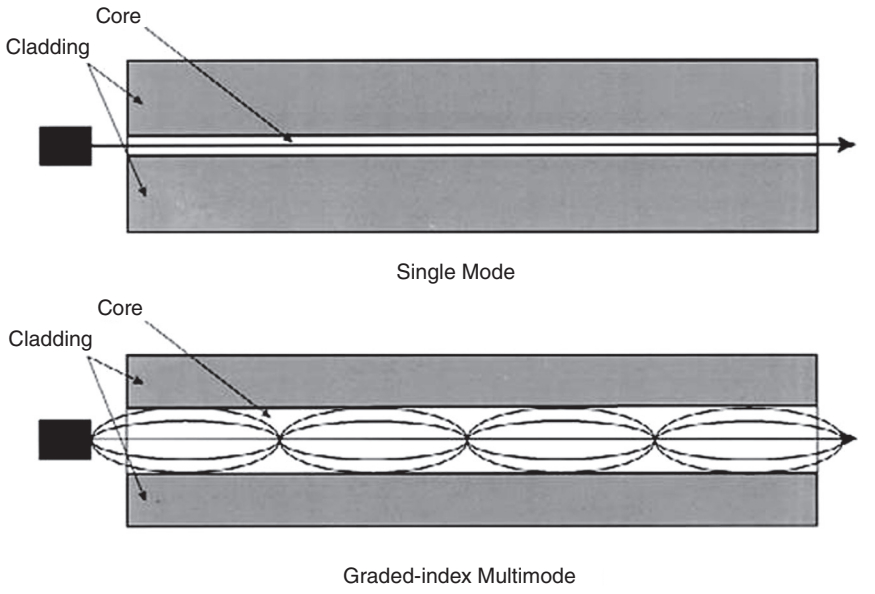


Exhibit 6.9 Optical Fiber Cable

(μm), or microns, in diameter and may be made from a variety of glass or plastic compounds. Surrounding the core is a layer called the *cladding*. The optical characteristics of the cladding are always different from the core's characteristics so that light signals traveling through the core at an angle will reflect back and stay in the core. The cladding may vary in thickness from a few to several hundred microns. The outermost layer is the *jacket*. Composed of plastic or rubber, the jacket's function is to provide the cable with physical protection from moisture, handling, and other environmental factors.

Two types of optical fiber cable are used for voice and data communications, differentiated by their transmission characteristics (see Exhibit 6.9). *Multimode fiber* (MMF) has a core diameter between 50 and 125 μm. Because this diameter is relatively large, light rays at different angles will be traveling through the core. This phenomenon, known as *modal dispersion,* has the effect of limiting the bit rate and/or distance of the cable. MMF cable is generally limited to a maximum cable length of 2 km. *Single-mode fiber* (SMF) eliminates the multiple path problem of MMF by using a thin core with a diameter of 2 to 8 μm. This thin-core cable results in a single propagation path so that very high bandwidths over large distances (up to 10 km) can be achieved.

SMF is the most expensive type of fiber and is usually used for long-haul data and telecommunications networks. MMF is commonly used on LANs; it is less expensive but can still handle the required data rates and distances.

Optical cable is extremely difficult to wiretap, but it's easy to cut. Such cables should be protected by shielded conduits, not placed in accessible locations such as next to a baseboard on the floor of a public corridor—in a hospital!

6.3.4 Wireless Media. Wireless LANs use radio signals to interconnect LAN nodes. Wireless LANs are increasingly common in environments where:

- It is difficult to install new wiring (e.g., in a building with asbestos in the walls).
- There are mobile users (e.g., in a hospital or car rental agency).

- Right-of-ways for wiring are hard to obtain (e.g., campus environments that span roadways).
- A temporary network is necessary (e.g., at a conference or meeting).
- Residential areas have no other networking facilities.
- Conference centers, hotels, and colleges and universities need wide and easy network access.

Wireless LANs generally employ infrared, spread spectrum, or microwave communications technology. *Infrared* (IR) is used for a variety of communications, monitoring, and control applications. It is also used for such non-LAN applications as home entertainment remote control, building security intrusion and motion detectors, medical diagnostic equipment, and missile guidance systems. For wireless LANs, the most common IR communications band uses signals with a wavelength in the range 800 to 1,000 nanometers (nm, or 10^{-9} m). *Diffused IR* operates at data rates between 1 to 4 Mbps at distances up to 200 feet, and can be used for stationary or mobile LAN nodes. *Directed Beam IR*, which requires line-of-sight, operates at data rates from 1 to 10 Mbps at distances up to 80 feet. IR systems are limited to a single room because the signals cannot pass through walls.

Spread spectrum is a wireless communications technology in the region of 2.4 or 5 gigahertz (GHz, or billions of cycles per second), where the actual frequency of the transmitted signal is deliberately varied during transmission. Originally, the frequency shifting was for security purposes to prevent monitoring of the communications channels. Two types of spread spectrum technology are used in LANs:

1. In *frequency hopping spread spectrum* (FHSS),[3] the transmitter sends the signal over a set of radio frequencies, hopping from frequency to frequency at split-second intervals in what appears to be a random sequence. The sequence is not random, however, and the receiver changes frequencies in synchronization with the transmitter. FHSS can support data rates from 1 to 3 Mbps up to a distance of 330 feet (100 m).
2. In *direct sequence spread spectrum* (DSSS), each bit in the original data stream is represented by multiple bits in the transmitted signal, spreading the signal across a wide frequency range. One result of DSSS is that the system can achieve a greater bandwidth than the original signal. DSSS can support data rates in excess of 20 Mbps up to a distance of 1,000 feet (300 m).

Another modulation scheme used in wireless LANs is *orthogonal frequency-division multiplexing* (OFDM). OFDM is a variant of frequency division multiplexing and uses a forward error-correction scheme and orthogonal subcarriers in order to minimize frequency crosstalk and bit errors.

Wireless access points can be purchased for less than $100, making this an attractive alternative to even UTP-based LANs in many scenarios.

Microwave LANs refers to communications in the area of 1, 5, and 19 GHz. Electromagnetic energy with a frequency higher than 1 GHz and data rates up to 20 Mbps can be maintained for distances up to 130 feet. One major disadvantage of microwave is that Federal Communications Commission (FCC) licensing is required for many of these frequencies.

6.3.5 Summary. In the early 1980s, coaxial cable was the most commonly used LAN medium. Twisted pair, used for telephony applications, was not used in LANs because high speeds could not be achieved. Optical fiber technology was still in its infancy and was very expensive. All of this changed by the early 1990s, when the electronics to drive twisted pair had dramatically improved, and optical fiber technology had greatly matured. It is rare to see coaxial cable used in a LAN today; instead, UTP (less costly than coax) or optical fiber (higher speeds than coax) are more often employed. Wireless LANs are a viable alternative to wire-based networks, yet more difficult to secure.[4] However, growth in wireless network access to the Internet outstripped fixed broadband subscriptions by the late 2000s; in 2011, rates of growth for wireless subscriptions were 200 to 300 percent of the rates of growth for fixed subscriptions.[5]

6.4 MEDIA ACCESS CONTROL. As mentioned, LANs are broadcast networks connecting peer devices, all having equal access to the medium. These characteristics place two requirements on the protocol that controls access to the network:

1. There can be only one station transmitting at any given time since multiple transmitters would result in garbled messages.
2. All stations must follow the same rules for accessing the network since there is no master station.

The schemes controlling access to the network medium are called *media access control* (MAC) protocols. Although many different LAN MAC schemes have been introduced in working products, the most common ones are essentially variants of two approaches: *contention* and *distributed polling*. These schemes will be discussed below, along with reference to appropriate Institute for Electronics and Electrical Engineers (IEEE) LAN standards.

6.4.1 Contention. A contention network can be compared to a group of people sitting around a conference table without a chairperson. When someone wants to speak, it is necessary first to determine whether anyone else is already speaking; if someone else is speaking, no one else can begin until that person has stopped. When a person detects silence at the table, he or she starts to talk. If two people start to talk at the same time, a *collision* has occurred and must be resolved. In the human analogy, collisions are resolved in one of two ways: Either both speakers stop and defer to each other ("polite backoff") or both continue speaking louder and louder until one gives up (a "rudeness algorithm").

The contention scheme used in LANs is similar to the polite backoff situation, and is called *carrier sense multiple access with collision detection* (CSMA/CD). CSMA/CD is one of the oldest LAN MAC schemes in use today, used originally in Ethernet and becoming the basis of the IEEE 802.3 standard (to be described). Although there have been other contention schemes used on LANs, CSMA/CD is the one that has survived and thrived in the marketplace.

CSMA/CD works on logical bus networks. When a station is ready to transmit, it first listens to the network medium ("carrier sense"). If the station detects a transmission on the line, it will continue to monitor the channel until it is idle. Once silence is detected, the station with a message to send will start to transmit. Stations continue to monitor the channel during transmission so that if a collision is detected, all transmitters stop transmitting.

CSMA/CD networks employ a *backoff scheme* so that the first collision does not bring the network down. Without a backoff scheme, all transmitters would detect a collision and stop transmitting; after again hearing silence on the line, however, all stations would once again start transmitting and would again collide with each other. The backoff scheme causes stations to make a random decision whether to transmit or not after silence is detected on the channel after a collision has occurred.

CSMA/CD uses a backoff scheme called *truncated binary exponential backoff*. Although this name is a mouthful, it actually describes the process very precisely. When a station is ready to transmit and detects silence on the line, it will attempt to send a message with a probability of 1 (i.e., 100 percent likelihood that it will transmit); this probability is called the *persistency* of the MAC scheme.[6] If a collision occurs, the station will stop transmitting and again wait for silence on the line. When silence is again detected, the station will transmit with a probability of $^1/_2$ (i.e., there is a 50 percent chance that it will transmit and a 50 percent chance that it will not). If two stations were involved in the collision and they both back off to a $^1/_2$-persistent condition, then there is a 50 percent chance that one will transmit and one will defer at the next transmission opportunity, a 25 percent chance that both will defer at the next opportunity, and a 25 percent chance that both will collide again.

If a station collides again, its persistency is again cut in half, now to $^1/_4$. All stations involved in the collision(s) drop their persistency and each station independently determines whether it will transmit at the next occurrence of silence or not.

As long as collisions occur, the persistency is continually cut in half until the station either successfully transmits or has 16 unsuccessful attempts to transmit the message. After 16 failed attempts, the station gives up.[7] After the station successfully transmits *or* has 16 unsuccessful attempts, the station's persistency returns to 1 and the operation continues as before.

Wireless LANs also use a form of contention, but it is generally not CSMA/CD because collision detection is not practical in a wireless environment. Instead, the stations still employ CSMA—they listen for an idle channel—but they do not necessarily transmit when the channel is idle. Instead, they wait to see if the channel remains idle for some period of time in an attempt to stave off a collision. This is a form of *CSMA with collision avoidance* (CSMA/CA).

6.4.2 Distributed Polling. Imagine that the same group of people is sitting around the same conference table, still without a chairperson. One person at the table has a microphone and can say anything to anyone in the room. Everyone in the room, of course, will hear the message. The rule here is that the only person who is allowed to speak is the one with the microphone; furthermore, the person will hold on to the microphone only while he or she has something to say and can hold on to it only for some maximum amount of time. When the first person is done talking or the time limit is reached, the microphone is passed to the next person at the table. Person 2 can now speak or immediately pass the microphone on to person 3. Eventually, the first person at the table will get the microphone back and get another opportunity to talk.

The scheme just described is implemented in LANs with a scheme called *token passing*. This is the basis for the IBM Token Ring and represents the second most commonly used LAN MAC algorithm. Token passing, in one variant or another, is the basis for the IEEE 802.4 and 802.5 standards, as well as for the Fiber Distributed Data Interface (FDDI).

Token passing requires a logical ring topology. When a station has data to send to another station, it must wait to receive a bit pattern representing the *token*. Tokens are

sent in such a way that only one station will see it at any given time; in this way, if a station sees the token, it has temporary, exclusive ownership of the network.

If a station receives the token and has no data to send, it passes the token on. If it does have data to send, it generates a *frame* containing the data. After sending the frame, the station will generate and send another token.

A *token ring* network is a logical ring implemented on a physical topology that supports a serial broadcast operation (i.e., a star or a ring). Each station receives transmissions one bit at a time and regenerates the bits for the next station. A station transmitting a frame will send the bits on its output link and receive them back on its input link. The transmitter, then, is responsible for removing its message from the network. When finished transmitting, the station transfers control to another station by sending the bits comprising a token on its output link. The next station on the ring that wants to transmit *and* sees the token can then send its data frame. Token rings (standardized in 802.5 and FDDI) are the most common implementation of token passing.

A *token bus* network (as specified in 802.4) is conceptually similar to the token ring, except that it is implemented using a simultaneous broadcast topology (i.e., a bus). In this physical topology, all stations hear all transmissions. A station that wants to send data to another will address a frame to the intended receiver on the network, as in a CSMA/CD bus. When done transmitting, the station will address a token to the next station logically in the ring; while all stations will hear the token transmission, only the one station to which it is addressed will pick it up. After receiving a token, a station may or may not transmit data, but it is, in any case, responsible for passing the token to the next station in the logical ring. Eventually, the token will return to the first station.

6.5 LAN PROTOCOLS AND STANDARDS.

The Open Systems Interconnection (OSI) Reference Model continues to be the standard framework with which to describe data communications architectures, including those for LANs. The basic LAN protocol architecture maps easily to the OSI model, as discussed in this section.

6.5.1 OSI Model versus LAN Model Architectures.

Although the LAN protocol architecture can be related to the OSI model, there is not a perfect one-to-one mapping of the protocol layers (see Exhibit 6.10). The OSI Physical Layer is analogous to a LAN Physical Layer (PHY). Both specify such things as:

- Electrical characteristics of the interface
- Mechanical characteristics of the connector and medium
- Interface circuits and their functions
- Properties of the medium
- Signaling speed
- Signaling method

Most LAN physical layer specifications actually comprise two sublayers. The lower sublayer describes physical layer aspects that are specific to a given medium; the higher sublayer describes those aspects that are media-independent.

The OSI Data Link Layer, responsible for error-free communication between any two communicating devices, is represented by two sublayers in a LAN. The lower sublayer is the MAC, which deals with issues of how the station should access the

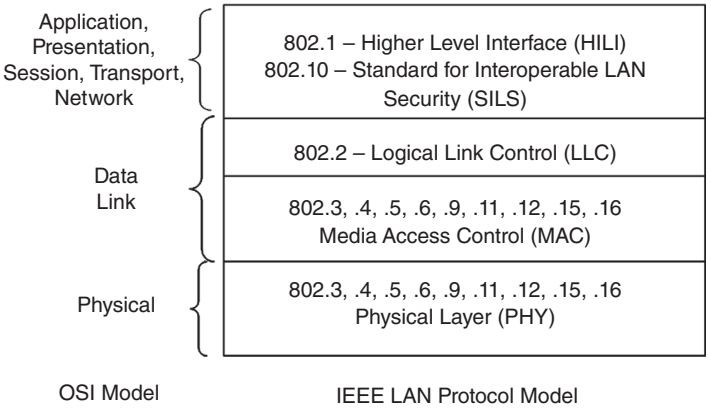


EXHIBIT 6.10 IEEE versus LAN Protocol Models

network medium. The MAC is responsible for error-free communication over the PHY and specifies such things as:

- Framing
- Addressing
- Bit-error detection
- Control and maintenance of the MAC protocol
- Rules governing medium access

The upper sublayer is called the Logical Link Control (LLC). The LLC protocol is responsible for maintaining a logical connection between two communicating LAN stations. The LLC specifies such rules as:

- Frame sequencing
- Error control
- Establishment and termination of a logical connection
- Addressing of higher layer services

Recalling that the main functions of the network layer are routing and congestion control, there are two reasons that no LAN protocol layer acts strictly like the OSI Network Layer:

1. There is no need for a routing algorithm in a broadcast network because all stations receive all transmissions; the address of the intended receiver is included in the transmission itself.
2. Congestion control is also not an issue in a broadcast network; a broadcast network must be limited to a single transmitter at a time, and this is accomplished by the MAC layer.

There are no standards for LANs corresponding to the upper four layers of the OSI model. Even in the less organized 1980s, end-to-end protocols as such were not required in a LAN environment because the end-to-end communication was limited to nodes on the LAN, and for that the MAC guaranteed error-free communication.

Only when LAN interconnection, via WAN and LAN access to the Internet, gained popularity did other end-to-end protocols become necessary. IP (and other network layer protocols) grew in demand as well. Those protocols are associated with the communications software as part of a network operating system (NOS), and these will be discussed later.

6.5.2 IEEE 802 Standards. Although they are not directly related to security, it is useful to be familiar with the standards describing LANs, the most common of which are the IEEE 802 standards. The IEEE Computer Society formed the Project 802 Committee in February 1980 to create standards for LANs as part of its more general work on standards for microprocessors; no other organization was making any similar standardization efforts. Originally, there was to be a single LAN standard, operating at a speed between 1 and 20 Mbps. The standard was divided into three parts: PHY, MAC, and a high-level interface (HILI) to allow other protocol suites to have a common protocol boundary with the LAN. The original MAC was based on the Ethernet standard, but other MAC schemes were quickly added and, over the years, the 802 committee has addressed many LAN schemes. They all have in common an interface to a single LLC protocol that provides a common interface between the HILI and any MAC.

A description of the Project 802 working groups (WG) and their status as of October 2012 follows.[8]

802.1—High-Layer LAN Protocols Working Group. Provides the framework for higher-layer issues, including protocol architecture, security, end-to-end protocols, bridging, internetworking, network management, and performance measurement.

802.2—Logical Link Control Working Group. Provides a consistent interface between any LAN MAC and higher-layer protocols. Depending on the options employed, the LLC can provide error detection and correction, sequential delivery, and multiprotocol encapsulation. The 802.2 standard is described in more detail in Section 6.5.6. This WG has been *disbanded*.[9]

802.3—Ethernet Working Group. Defines the MAC and PHY specifications for a CSMA/CD bus network. This specification is discussed in more detail in Section 6.5.3. (The 802.3 CSMA/CD standard is based on Ethernet, described in Section 6.5.4.)

802.4—Token Bus Working Group. Defines the MAC and PHY specifications for a token-passing bus based on work originally done at General Motors as part of the Manufacturing Automation Protocol (MAP). Well suited for factory floors and assembly lines, MAP never achieved widespread use. This WG has been disbanded.

802.5—Token Ring Working Group. Defines the MAC and PHY specifications for a token-passing ring. Although this WG is currently disbanded, the specification is discussed in more detail in Section 6.5.5 for historical purposes.

802.6—Metropolitan Area Network (MAN) Working Group. Defines the MAC and PHY specifications for a MAN. In particular, the 802.6 standard defines a MAC and PHY called Distributed Queue Dual Bus (DQDB), which was one of the MACs employed with the Switched Multimegabit Data Service (SMDS) and Connectionless Broadband Data Service (CBDS). Introduced in the early 1990s, neither service remains in use today. This WG has been disbanded.

802.7—Broadband Technology Advisory Group (BBTAG). Advises other 802 subcommittees about changes in broadband technology and their effect on the 802 standards. This WG has been disbanded.

802.8—Fiber Optics Technology Advisory Group (FOTAG). Advises other 802 subcommittees about changes in optical fiber technology and their effect on the 802 standards. This WG has been disbanded.

802.9—Integrated Services LAN (ISLAN) Working Group. Defines the MAC and PHY specifications for integrated voice/data terminal access to integrated services networks, including ISLANs and MANs, and Integrated Services Digital Networks (ISDN). The only practical implementation was deployed in IsoEthernet products, described in the IEEE 802.9a standard. This WG has been disbanded.

802.10—Security Working Group. Defines procedures for providing security mechanisms on interconnected LANs, including cryptography and certificates. This WG has been disbanded.

802.11—Wireless LAN (WLAN) Working Group. Defines MAC and PHY specifications for "through the air" media. The original 802.11 standard defined operation at 1 or 2 Mbps using the 2.4-GHz range and DSSS or FHSS spread spectrum technology; nominal maximum distances were 330 feet (100 m). The most common variants today are:

- 802.11b—Data rates up to 11 Mbps at a nominal maximum distance up to 460 feet (140 m) on a frequency of 2.4 GHz using DSSS.
- 802.11g—Data rates up to 54 Mbps at a nominal maximum distance up to 460 feet (140 m) on a frequency of 2.4 GHz using DSSS or OFDM.
- 802.11n—Data rates up to 150 Mbps at a nominal maximum distance up to 820 feet (250 m) on a frequency of 2.4 or 5 GHz using OFDM.

Future 802.11 standards are expected that will provide data rates up to 866.7 Mbps on a frequency of 5 GHz using OFDM.

802.12—Demand Priority Working Group. Describes one of the MAC and PHY specifications originally proposed for 100 Mbps LAN speeds and dubbed *100BASE-VG/AnyLAN*. Largely unused, and the WG has been disbanded.

802.13. (This number was never assigned to a WG because it was felt that the *13* would hamper products in the marketplace.)

802.14—Cable Modem Working Group. Originally intended to describe LANs for cable TV systems. This WG has been disbanded.

802.15—Wireless Personal Area Network (WPAN) Working Group. Defines a MAC and PHY for a short distance wireless network between portable and mobile devices such as PCs, personal digital assistants (PDAs), cell phones, pagers, and other communications equipment.

802.16—Broadband Wireless Access (BBWA) Working Group. Defines the MAC and PHY for high-speed wireless network access over relatively short distances. BBWA standards address the "first-mile/last-mile" connection in wireless metropolitan area networks, extending the reach of residential broadband services such as cable modem or digital subscriber line (DSL).

802.17—Resilient Packet Ring (RPR) Working Group. Defines standards to support the development and deployment of RPR local, metropolitan, and wide

area networks for resilient and efficient transfer of data packets at rates scalable to many gigabits per second. This WG is currently in *hibernation*.[10]

802.18—Radio Regulatory Technical Advisory Group (RR-TAG). On behalf of other 802 WGs using radio-based communication, this TAG monitors, and actively participates in, ongoing national and international radio regulatory activities.

802.19—Wireless Coexistence Working Group. Develops and maintains policies defining the responsibilities of 802 standards developers to address issues of coexistence with existing standards and other standards under development.

802.20—Mobile Broadband Wireless Access (MBWA) Working Group. Defines the specification for a packet-based wireless interface that is optimized for IP-based services. The goal is to enable worldwide deployment of affordable, ubiquitous, always on, and interoperable multivendor mobile broadband wireless access networks that meet the needs of business and residential end user markets. This WG is currently in hibernation.

802.21—Media Independent Handover Services Working Group. Developing standards to enable handover and interoperability between heterogeneous network types including both 802 and non-802 networks.

802.22—Wireless Regional Area Networks (WRAN) Working Group. Developing a standard for a radio-based PHY, MAC, and air interface for use by license-exempt devices on a noninterfering basis in the spectrum allocated to broadcast television.

802.23—Emergency Services Working Group. This working group was created to define an IEEE 802 framework for LANs that would comply with applicable civil authority requirements for communications systems. This working group has been disbanded.

802.24—Smart Grid Technology Advisory Group. This TAG was created to provide liaison between the 802 committee and the smart grid industry and regulatory bodies, and to provide coordination and collaboration amongst 802 working groups related to smart grids.

6.5.3 IEEE 802.3 CSMA/CD Standard.

The original IEEE 802.3 standard, first published in 1985, describes the PHY and MAC for a CSMA/CD bus network operating over thick coaxial cable. Today, an 802.3 network implementation can employ any of a number of media types, including UTP and optical fiber. Without question, star-wired UTP implementations are the most popular.

The 802.3 committee anticipated the different media types that might be used, and they developed a nomenclature to identify the actual physical implementation, using the format:

[speed (Mbps)][signaling type][segment length (m) or media type]

The original 802.3 specification, for example, operated at 10 Mbps, used baseband (digital) signaling and limited a single coaxial cable segment to a length of 500 m (1,640 feet); the cable was designated 10BASE5. In fact, the largest distance between two 802.3 stations could be 2.8 km (9,200 feet), so repeaters might be used to interconnect several 500-m coaxial cable segments.

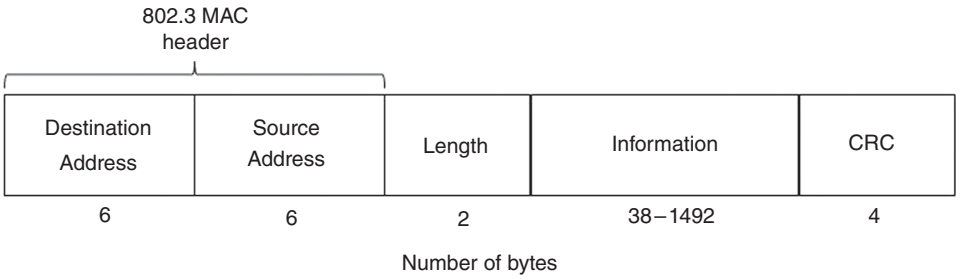


EXHIBIT 6.11 IEEE 802.3 Frame Format

A less expensive version, called *CheaperNet*, was later introduced that operated over thin coaxial cable segments limited to 185 m (610 feet); this PHY is denoted 10BASE2.

In the mid-1980s, AT&T introduced a product called StarLAN, which operated at 1 Mbps over UTP. Although this product has long been relegated to obscurity, it was the first to break the 1-Mbps barrier on UTP. Subsequent versions of 802.3 that employ UTP all use a star topology where each network node connects directly back to a central hub. The first 10-Mbps version of 802.3 was denoted 10BASE-T, the *T* indicating use of the UTP medium (which structured wiring standards say is limited to a distance of 100 m, or 330 feet). The 10-Mbps optical fiber version of 802.3 is 10BASE-F. Today, of course, 100-Mbps and 1-Gbps versions (i.e., 100BASE-T and 1000BASE-T) are available. Full-duplex Ethernet takes advantage of the point-to-point links in a star configuration and effectively doubles the line speed by allowing both stations to transmit at the same time.

Exhibit 6.11 shows the format of an IEEE 802.3 MAC frame, primarily for reference purposes. The fields and their functions are:

- **Preamble**. Used for clock synchronization; employs 7 repetitions of the 8-bit pattern 10101010. (8 binary bits = 1 byte = 1 octet)
- **Start frame delimiter (SFD)**. The bit pattern 10101011 denotes the actual beginning of the frame. 1 octet.
- **Destination address (DA)**. 48-bit MAC address of the station that should receive this frame. An all-1s address in 48 binary bits (ff-ff-ff-ff-ff-ff in hexadecimal) is the *broadcast address*, indicating that all stations should receive this message.
- **Source address (SA)**. 48-bit MAC address of the station sending this frame.
- **Length**. Number of octets in the LLC data field, a value between 0 and 1500. 2 octets.
- **LLC Data**. Data from LLC (and higher layers). This field contains a 3-octet LLC Header, 5-octet 802.2 Subnetwork Access Protocol (SNAP) header, and 38 to 1492 octets of higher layer data.
- **PAD**. Additional octets to ensure that the frame is at least 64 octets in length; this minimum is required by CSMA/CD networks as part of the collision detection mechanism.
- **Frame check sequence (FCS)**. Remainder from CRC-32 calculation used for bit error detection. 4 octets.

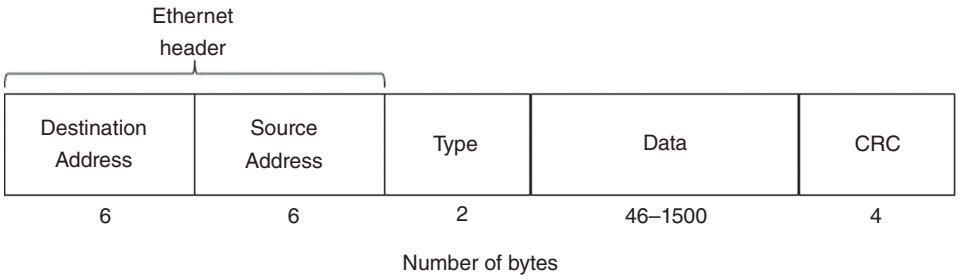


EXHIBIT 6.12 Ethernet II Frame Format

6.5.4 Ethernet II. The IEEE's CSMA/CD standard is based on the Ethernet specification developed at Xerox's Palo Alto Research Center (PARC) in the mid-1970s. When Xerox first decided to market Ethernet, there was no OSI model or any LAN standards or products. Given that environment, Xerox sought industry support for this new specification. The Ethernet specification has been jointly distributed (and marketed) by Digital Equipment Corporation (DEC, now Compaq), Intel, and Xerox (hence sometimes known as *DIX Ethernet*). While the 802.3 standard is based on Ethernet II, the two are not exactly the same.

Exhibit 6.12 shows the format of an Ethernet MAC frame, primarily for purposes of comparison to the IEEE frame. The fields and their functions are:

- **Preamble**. Used for clock synchronization; employs the bit pattern 10101010 … 10101011. 8 octets.
- **Destination address (DA)**. 48-bit MAC address of the station that should receive this frame. An all-1s address (ff-ff-ff-ff-ff-ff) is the *broadcast address*, indicating that all stations should receive this message.
- **Source address (SA)**. 48-bit MAC address of the station sending this frame.
- **Protocol identifier (PID)**. Indicator of the protocol information transported in the Information field. Sample values include 2048 and 2054 to indicate the Internet Protocol (IP) and Address Resolution Protocol (ARP), respectively. 2 octets.
- **Information**. Protocol data unit from the protocol identified in the PID field. 46 to 1,500 octets. (It is the responsibility of the higher layer to ensure that there are at least 46 octets of data in the frame.)
- **Frame check sequence (FCS)**. Remainder from CRC-32 calculation used for bit error detection. 4 octets.

The point in comparing the frame formats of Ethernet and 802.3 is primarily of historical purposes because today's implementations are 802.3 and not Ethernet. That said, it is interesting to note that the two specifications are, in fact, different. It is a minor thing, perhaps, and was a common misnomer in the industry to refer to *IEEE 802.3 Ethernet* (even the IEEE 802.3 committee is now known as the *Ethernet Working Group*), but it was an important difference to both a network administrator and a security professional.

In particular, in years past, if one LAN device only understood Ethernet encapsulation, it would not be able to communicate successfully with another LAN device that only understood IEEE 802.3 encapsulation. Both devices, however, can share the same

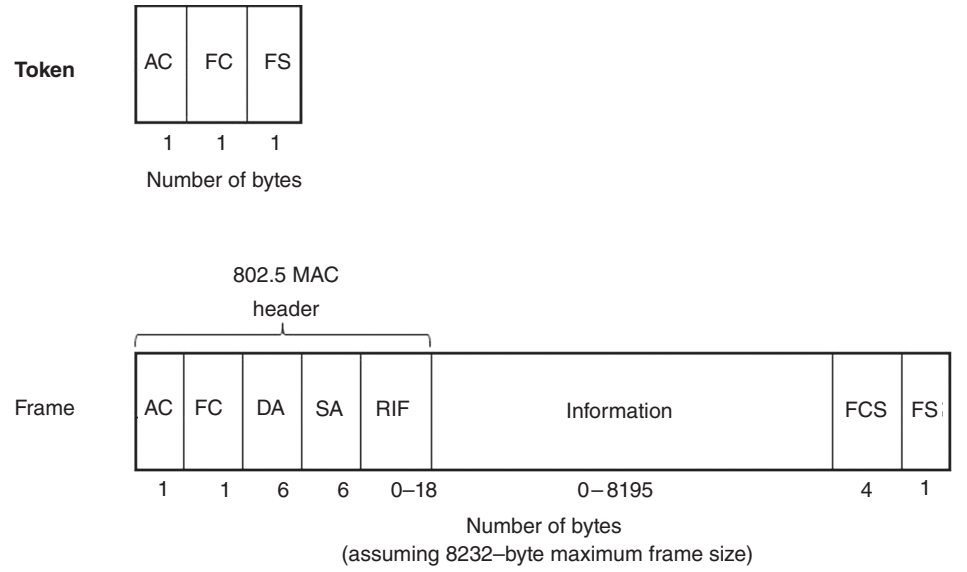


EXHIBIT 6.13 IEEE 802.5 Token and Frame Formats

medium backbone because the electronics are the same. A NetWare server running the Internetwork Packet Exchange (IPX) network layer protocol over IEEE 802.3 frames, for example, could easily share the network with a UNIX host running IP over Ethernet, but it maintained some immunity from attack by an IP host for the NetWare server because the two networks could not cross-communicate.

6.5.5 IEEE 802.5 Token-Ring Standard. The IEEE 802.5 token-ring standard was based on the IBM product of the same name. Both the standard and the product date back to about 1985. It is described here for historical purposes.

The token ring has a logical ring topology, although it was built as a physical star. Designed to operate with STP or UTP cable, most implementations operated at speeds of 16 Mbps or higher. The 802.5 MAC was essentially the same as the token passing scheme described in Section 6.4.2. The fields of the MAC frame (see Exhibit 6.13) are:

- **Start delimiter (SD)**. Marks the actual beginning of the transmission. Bit pattern JK0JK000, where J and K represent special symbols on the line.[11] 1 octet.
- **Access control (AC)**. Indicates whether this transmission is a *token* (i.e., no data) or a *frame* (i.e., contains data). This field also contains information about the priority of this transmission. 1 octet.
- **Frame control (FC)**. Indicates if this frame carries LLC (and higher-layer) data or MAC management information; if it is MAC-specific information, this field also indicates the MAC frame type. 1 octet.
- **Destination address (DA)**. 48-bit MAC address of the station to which this frame is intended.
- **Source address (SA)**. 48-bit MAC address of the station sending this frame.
- **Routing information field (RIF)**. An optional field, used only in multiple-ring networks utilizing source routing *and* in which the intended receiver is on a different ring than the transmitter. In source routing, the transmitter can specify

the intended path of this frame, designating up to eight intermediate networks.[12] 0 to 18 octets.

- **Information (INFO)**. Contains an LLC frame or MAC management information. No maximum length is specified by the standard, but the length of this field will be limited by the time required to transmit the entire frame, controlled by the *token holding time* parameter.
- **Frame check sequence (FCS)**. Remainder from a CRC-32 calculation to detect bit errors in the frame. 4 octets.
- **End delimiter (ED)**. Demarks the end of the transmission, with the bit pattern JK1JK1IE, where J and K are as described in the SD field. The I-bit indicates whether this frame is the last frame of a multiple-frame sequence and the E-bit indicates whether a bit error was detected by the receiver (E); these bits are cleared by the original sender when the frame returns to that station. 1 octet.
- **Frame status (FS)**. The bit pattern AC00AC00; these bits indicate whether the frame's destination address was recognized by any station on the network (A) and whether this frame was successfully copied by the intended receiver (C). 1 octet.

As shown, a token comprises just three octets, the SD, AC, and ED fields. A station sends a frame whenever there is user data or MAC information to send. The station must wait until it receives a *token* before it can generate a *frame*.

The transmitting station is responsible for generating a new token after it transmits a single frame. Recall that the transmitted bits come back to the sender, and it is this station that removes the bits from the network. According to the original standard, the transmitter will send a token after sending all of the bits of the frame and must wait until it has seen at least the returning SA field to verify that it is, in fact, removing its own frame from the network. Optionally, early token release allows the transmitter to generate a new token immediately after finishing sending the bits from its frame, even if the SA field has not yet returned. This latter option was developed to improve performance in very large token ring environments, such as the American National Standards Institute (ANSI) FDDI standard.

Today, 802.5 token rings are primarily limited to IBM environments, and there is a lot to be found there. FDDI is more commonly found in multibuilding campus environments, used as a backbone to interconnect Ethernet/802.3 networks. FDDI is being phased out; the last FDDI product vendor dropped out of the marketplace in 1999.

6.5.6 IEEE 802.2 LLC Standard. The IEEE 802.2 LLC protocol was intended to provide a common interface between 802 LAN MACs and higher-layer applications. With the LLC, the underlying MAC scheme is transparent to the application just as the application is transparent to the MAC.

The LLC was designed to support any number of services, the most common being an unacknowledged connectionless service (primarily used in contention networks) and an acknowledged connection-oriented service (primarily used in token ring environments).

The LLC is loosely based on the Higher-layer Data Link Control (HDLC) bit-oriented protocol in both operation and frame format (see Exhibit 6.14). The LLC frame appears in the Information field of a MAC frame. The first two fields of the LLC header are the *Destination Service Access Point* (DSAP) and the *Source Service Access Point* (SSAP) fields, originally intended to identify the higher-layer services at the source

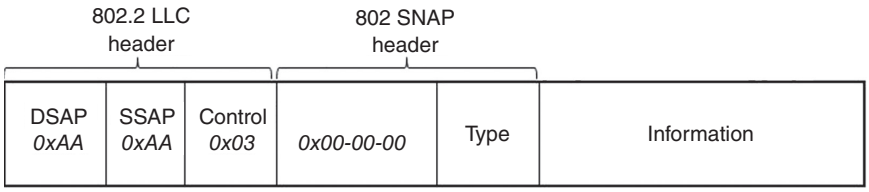


EXHIBIT 6.14 IEEE 802.2 LLC Frame Transporting SNAP Header (which in turn indicates IEEE organization and EtherType protocol identifiers)

and destination node. This is similar in concept to *ports* in the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) but was never well implemented, and the DSAP and SSAP values are typically the same. The third field is the Control field, identifying the type of frame.

The Subnetwork Access Protocol (SNAP) is an IEEE 802 protocol that can be used to identify *any* protocol created by *any* agency, and is commonly used above the LLC layer. In this case, the SNAP header immediately follows the LLC header. Use of SNAP is indicated by the LLC fields when both DSAP and SSAP fields are set to a value of 170 (0xAA) and the Control field is set to a value of 3 (octal 03) to indicate that it is an Unnumbered Information frame.

The SNAP header has two fields. The 3-byte *Organizationally Unique Identifier* (OUI) field refers to the organization that developed either the higher-layer protocol or a way to refer to the protocol. The 2-byte *Type* field identifies the protocol using the Organization-defined number.

The Internet Protocol (IP) and Address Resolution Protocol (ARP) provide example uses of SNAP. The common format of a SNAP header encapsulating these protocols would be to set the OUI value to 0 (0x00-00-00) to identify IEEE/ISO as the organization. The Type field would then use the *EtherType* values of 2048 (0x08-00) and 2054 (0x08-06) to indicate use of IP and ARP, respectively.

6.5.7 Summary. This section has covered the important LAN standards governing what is most likely to be seen in the industry today. Table 6.1 summarizes some of the discussion about the most common LAN topologies, media, MAC schemes, and standards.

6.6 INTERCONNECTION DEVICES. LAN interconnection devices are used to attach individual LANs to each other in order to build a large enterprise network. They can also interconnect LAN components across a WAN and provide LAN access to the

TABLE 6.1 LAN Characteristics

Physical Topology	Logical Topology	Media	MAC	Speed (Mbps)	Standard
Bus	Bus	Coax	CSMA/CD	10	802.3, Ethernet
Star	Bus	UTP, Fiber	CSMA/CD	1–1000	802.3
Star	Bus	Wireless	CSMA/CA	1–150	802.11
Star	Ring	UTP	Token passing	16	802.5
Ring	Ring	Fiber	Token passing	100	FDDI

Internet. Several types of such devices are used for LAN interconnections, including hubs, switches, bridges, and routers. The major distinction between these devices is the OSI layer at which they operate, and all are discussed in the next sections.

6.6.1 Hubs. *Hubs* are used to build physically star-wired LANs, using media that are basically point-to-point in nature (such as UTP and optical fiber). Note that it is the internal wiring of the hub that determines its *logical* nature, so that a logical bus or ring LAN can be physically star-wired.

So-called *Ethernet hubs* support 10, 100, and/or 1,000 Mbps Ethernet or 802.3 networks. Different hubs will have a different number of ports, generally ranging from 4 to 32. Hubs provide physical connectivity only; when a frame arrives on one port, the hub will broadcast the frame back out to all other ports, which simulates the broadcast bus environment. Multiple hubs can be interconnected to form reasonably large networks.

Token-ring hubs, generally called *multistation access units* (MAUs), look similar to Ethernet hubs but have different internal wiring. When an MAU receives a transmission on one port, it merely forwards that transmission, a bit at a time, to the next port sequentially on the MAU. In this way, it simulates the ring environment.

6.6.2 Switches. *Switches* are generally employed in the CSMA/CD environment and extend the capabilities of a hub. A switch operates at a combination of PHY and MAC layers. In addition to providing physical connectivity like a hub, a switch learns the MAC address of all stations attached to it. When a frame arrives on a switch port, the switch looks at the destination MAC address and places the frame on the port associated with that address (which might be the port leading to another switch).

Switches are used primarily to improve performance. Given the scenario described earlier, multiple stations can transmit simultaneously without collision. Furthermore, switches can operate in full-duplex mode, meaning that a single station can both transmit and receive at the same time. A 10 Mbps switched Ethernet LAN, for example, can achieve performance similar to that of a 100-Mbps hubbed Ethernet LAN. (This was a real boon in those environments where it is not viable to upgrade 10-Mbps NICs and wiring.)

There is a subtle security ramification to the use of switches versus hubs. In particular, if a user places a packet sniffer on a hubbed LAN, the sniffer will see every frame because the hub simulates the broadcast environment. A packet sniffer on a switched network will not be as effective; it will only pick up those frames that are specifically addressed to the LAN broadcast address. That said, many switches come with an administrative port that can be set to monitor all ports for troubleshooting purposes.

6.6.3 Bridges. A *bridge* provides a point-to-point link between the two LANs, usually those employing similar MAC schemes. Bridges operate at the MAC layer, and their operation is controlled by the MAC address.

Ethernet environments commonly employ *learning bridges*. In a very simple case, consider a bridge interconnecting two LANs, #1 and #2 (see Exhibit 6.15). When any LAN station sends a frame, both destination and source MAC addresses are included in the transmission. As frames appear on the networks, the bridge sees all of the source addresses and builds a table associating the MAC addresses with one LAN or the other, eventually learning the location of all of the network's stations. This process is sometimes called *backward learning* because the bridge learns the location of stations that transmit.

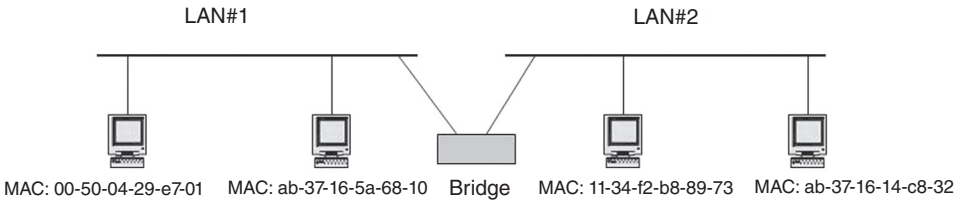


EXHIBIT 6.15 Two LANs Interconnected via a Bridge

A bridge is a simple frame store-and-forward device. Like all stations on the LAN, a bridge examines the destination address of any transmitted frames. If a transmission on LAN #1 contains a destination address of a station on LAN #2, the bridge will forward the frame. If a transmission contains an unknown destination address, the bridge will also forward the frame.

Although a bridge bases its decisions on the MAC address, it is not an intelligent device; that is, it knows that a station with a particular MAC address is in one direction or another, but it does not know precisely where that station is. Because bridges have to build tables containing all of the stations' addresses that they learn, bridges do not scale particularly well to large networks. Bridges also extend the *broadcast domain* (i.e., if a frame transmitted on LAN #1 is sent to the broadcast address, it will be forwarded to LAN #2).

6.6.4 Routers. A *router* is conceptually similar to a bridge in that it is also a store-and-forward device. A router, however, works at the Network Layer and is therefore a much more powerful device than a bridge. As Exhibit 6.16 shows, every LAN device has both a MAC (hardware) and Network Layer (software) address (in this case, IP is the sample Network Layer address). Because Network Layer addresses are hierarchical, the networks themselves have a network identifier number (the NET_ID in Exhibit 6.16). Network Layer addresses are well suited to environments where intermediate devices have to find a best route between networks.

Like a bridge, a router is considered to be just another station on a LAN to which it is attached. If the router sees a transmission on LAN #1 (with a NET_ID of 192.168.16.0) containing a destination address of a station on another network, it will route the packet to the correct destination network, even if that means going through another router to get there.

This example also demonstrates another major difference between bridges and routers. In a bridged environment, a station on LAN #1 sends a frame to some MAC address and has no knowledge of whether the intended destination is on the same LAN or not; the bridge will forward the frame if necessary, but this is all transparent to sender

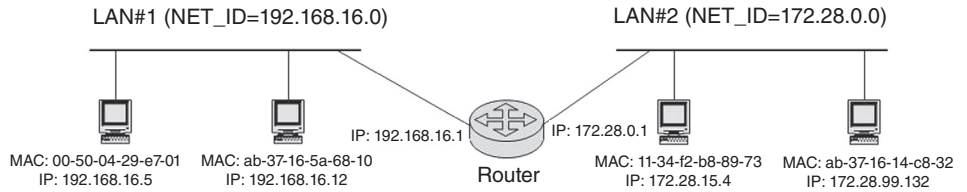


EXHIBIT 6.16 Two LANs Interconnected via a Router

and receiver. In a routed environment, however, the sender can tell if the receiver is on the same or different network merely by examining the destination network address. In fact, the router only gets involved if the packet has to leave the local network; that is why in an IP environment, for example, an address of a default gateway (router) has to be provided.

Routers also limit the broadcast domain. If a station on LAN #1 transmits a frame using the broadcast MAC address, the frame goes no further than the router.

Routers build their routing tables very differently than bridges. Whereas bridges learn the relative location of a station by observing a frame's source address, packets learn the Network Layer address by the use of routing protocols that allow groups of routers to exchange routing information.[13]

6.6.5 Summary. Hubs, switches, bridges, and routers are all commonly employed LAN interconnection devices. These are tools in the kit of everyone who works with LANs, as the building blocks of everything from small and intermediate-size local networks to large enterprise networks and the global Internet.

6.7 NETWORK OPERATING SYSTEMS.

Just as an operating system manages computer resources, a *network operating system* (NOS) provides the software that controls the resources of a LAN. NOSs generally comprise software that provides at least these functions:

- *Hardware drivers* are the software that allows the NOS to communicate with the NIC.
- *Communications software* allows applications running on different LAN nodes to communicate.
- *Services* are the functional aspects of the NOS and the reason that people use a LAN in the first place. Sample services include file services (file sharing), print services (commonly shared printers), message services (email), communication services (LAN access to the Internet), and fax services (commonly shared facsimile).

NOSs are typically classified as being peer to peer or client/server. A *peer-to-peer* LAN allows any LAN node to communicate with any other LAN node, and any LAN node can provide services to other nodes. In a *client/server* (or *server-based*) environment, every node is either a client or a server. In this scenario, *servers* are special nodes that offer services to other servers or to clients, while *clients* are the ordinary end-user workstations. Clients can only communicate with a server.

When evaluating or investigating the security of a LAN, the software is the most common point of exposure, vulnerability, and exploitation, particularly for remote attacks.

Some sample NOSs that have had historical significance include:

AppleTalk. Apple Macs have come with integrated LAN capabilities since their inception in 1985. Originally using a scheme called LocalTalk, AppleTalk was a peer-to-peer network running over a 10-Mbps CSMA/CD LAN. The Network Layer protocol historically associated with AppleTalk was called the

Datagram Delivery Protocol (DDP). Apple dropped support for AppleTalk in 2009 and supports TCP/IP-based networking.

Microsoft Networking. Microsoft operating systems have come with LAN capabilities since Windows for Workgroups (WfW or Windows 3.11). Employing a nonroutable protocol called the Network Basic Input/Output System (NetBIOS) Extended User Interface (NetBEUI), Windows client systems (Windows 3.11 and later) can be easily used to build an inexpensive, simple peer-to-peer LAN for file and print sharing. NetBEUI is nonroutable because it does not provide an addressing mechanism to allow interconnected yet distinct NetBEUI subnetworks; if two NetBEUI networks are attached in any way, they will appear to be one large network. (This is why the hard drive of improperly configured Windows systems can be viewed from across the Internet.) Microsoft defined NetBIOS over TCP/IP (NBT) in the 1980s to support encapsulation of NetBIOS messages in TCP and UDP messages. Most Microsoft networking today relies more on TCP/IP than NetBEUI.

Microsoft Windows Server (including Windows NT Server, Windows Server 2003, Windows Server 2008, and Windows Server 2012). The NOS of choice for Windows network environments, Windows Server is Microsoft's client/server operating system. Client systems can run nearly any Windows operating system, from Windows XP or NT Workstation to Windows VISTA, Windows 7, or Windows 8. Clients on a Windows Server network themselves can form a peer-to-peer network. Microsoft networking is reliant on Active Directory and TCP/IP.

Novell NetWare. Novell offered one of the first PC-class networks in the early 1980s using a proprietary star-based LAN. By the 1990s, NetWare was the best-known client/server NOS and accounted for more than 70 percent of the NOS market. The Network Layer protocol associated with classical NetWare is the Internetwork Packet Exchange (IPX) protocol. By the turn of the century, Microsoft was eating dramatically into NetWare's popularity and TCP/IP was dominating as the protocol-of-choice. At around this time, Novell embraced Linux and TCP/IP, and NetWare has been replaced with the Open Enterprise Server.

UNIX and Linux. TCP/IP has been the network communications protocol for UNIX systems since 1984, allowing UNIX-based hosts to build client/server (and, ultimately, peer-to-peer) networks. TCP/IP has also been integral to Linux since its inception in 1991. With TCP/IP, any system can run server (daemon) software to provide services to other systems, so that any system can act as a client or server, depending on application.

6.8 SUMMARY. Exhibit 6.17 shows a possible network design that includes many of the elements that have been described in this chapter (and a few that have not). This network's router provides the interface to the Internet and is attached via some sort of dedicated connection, such as a point-to-point 56 Kbps or T1 (1.544 Mbps) leased line, frame relay, or digital subscriber line.

In this scenario, the router is physically located at the main site. From a security perspective, the organization may segment its network into an *external* and *internal* side, the internal being protected by a firewall.[14] The external network includes the router, public Web server, and firewall. Those three systems are interconnected through

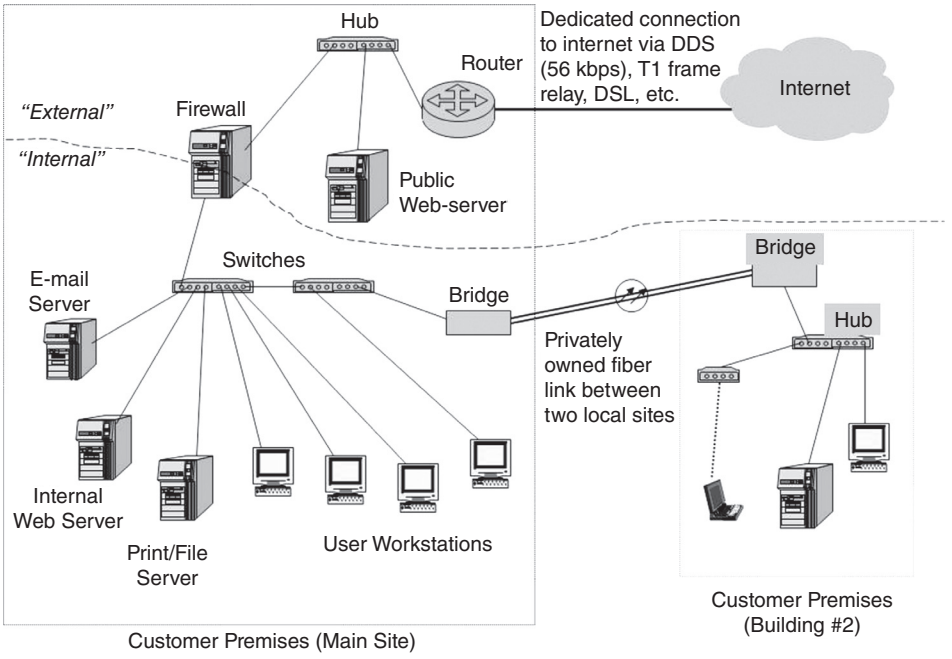


EXHIBIT 6.17 LAN Scenario

a hub to which they each attach via a Cat 5 UTP cable. In this scenario, the hub could actually implement 10BASE-T or 100BASE-T Ethernet, or even a token ring.

The external and internal networks are connected through the firewall, which, in this case, will have two NICs. The two networks are separate and distinct; the firewall does not extend the broadcast domain of either network, and, in fact, these two networks would have different IP network identifiers.

The internal network at the main site is a collection of servers and user workstations that are interconnected via a set of switches. In this example, these are 8-port 100-Mbps Ethernet switches. Since there are more than 8 devices, the switches themselves need to be interconnected. There are several options for that:

- *Stackable* switches physically attach to each other, extending the switch's backplane to create a larger switch (in this case, a 16-port switch).
- An *optical fiber* link can be used to interconnect the switch, usually at backplane speeds in the 1+-Gbps range.
- A *UTP* link might be used to interconnect the switches via two of the 100-Mbps ports.

To connect the LAN in Building #2 with the LAN at the main site, a point-to-point connection between a pair of bridges would suffice. In this case, the buildings are several kilometers apart, necessitating use of optical fiber.

In Building #2, there is another hub-based LAN, with a laptop using wireless technology, communicating with an access node that is also attached to the hub.

This chapter has only skimmed the surface of LAN concepts, standards, and technologies. Their study is important to the security professional, however, because LANs are the basis of all networking. As a *network of networks*, the Internet comprises

millions of local networks. This chapter indicates many of the points of potential vulnerability or compromise in a system.

6.9 FURTHER READING

Cisco Systems. *Internetworking Technology Handbook*. Cisco DocWiki Website. June 20, 2012. http://docwiki.cisco.com/wiki/Internetworking_Technology_Handbook

Gast, M. "Wireless LAN Security: A Short History." O'Reilly | Wireless Devcenter Website. April 19, 2002. www.oreillynet.com/pub/a/wireless/2002/04/19/security.html

Mikalsen, A., and P. Borgesen. *Local Area Network Management, Design and Security: A Practical Approach*. Hoboken, NJ: John Wiley & Sons, 2002.

Riley, S., and R. A. Breyer. *Switched, Fast, and Gigabit Ethernet*, 3rd ed. Indianapolis, IN: New Riders Publishing, 1998.

Stallings, W., and T. Case. *Business Data Communications—Infrastructure, Networking and Security*, 7th ed. Upper Saddle River, NJ: Prentice-Hall, 2012.

6.10 NOTES

1. See Chapter 5 in this *Handbook* for additional background information on LANs and WANs.
2. See Sections 22.4.5 and 23.9.1 in this *Handbook* for further discussion of wiretaps.
3. It is an interesting point of trivia to note that the original frequency-hopping spread spectrum technique was coinvented by Hollywood star Hedy Lamarr in 1940 and given, free of charge, to the U.S. Navy.
4. See Chapter 33 in this *Handbook* for more details of wireless LAN security.
5. M. Kende, "Internet Global Growth: Lessons for the Future," Analysys Mason Knowledge Centre (Website), 2012, www.analysysmason.com/internet-global-growth-lessons-for-the-future (p. 14).
6. Since CSMA/CD transmits with a probability of 1, it is sometimes referred to as being 1-persistent.
7. As an aside, although the station can experience 16 collisions, the probability of transmission will never fall below 1/1024, or 2-10, since Ethernet and IEEE 802.3 do not allow more than 1,024 devices on the network. This is the source of the word "truncated" in the name of the scheme.
8. Up-to-date status information about the 802 committee can be found at the LAN/MAN Standards Committee Web site at http://grouper.ieee.org/groups/802/.
9. A WG is disbanded when it is considered that there is no more work for the IEEE to undertake in this topic area.
10. A WG will go into hibernation when there are no new projects to undertake. This status indicates a WG that has reached status quo.
11. The term "special symbol" requires explanation. The signaling scheme used in the token ring PHY standard is called Differential Manchester. In this signaling scheme, the signal is at a positive voltage for half of the bit time and at a negative voltage for the other half of the bit time, meaning that each bit has a sum total of 0 volts (resulting in what is sometimes called DC balancing). The J and K symbols are Differential Manchester code violations, where one symbol is at negative voltage for an entire bit time and the other at positive voltage for an entire bit

time. These code violations have the benefit of being able to indicate special events and can be used for synchronization. J and K symbols are always used in pairs to maintain DC balancing.

12. Source routing is a very rarely used option in IP and is, in fact, a security problem; firewall administrators routinely set up filters to block IP packets with source routing. Source routing in an 802.5 network, however, is a normal feature and is not considered to be a security threat because this information has no impact on the WAN.
13. In the IP environment, common routing protocols include the Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and Routing Information Protocol (RIP).
14. This is a very simplistic firewall design with the internal and external network. The focus of this diagram is on the LAN components, however, rather than the specific security architecture.

CHAPTER 7

ENCRYPTION

Stephen Cobb and Corinne LeFrançois

7.1 INTRODUCTION TO CRYPTOGRAPHY 7·1
- 7.1.1 Terminology 7·2
- 7.1.2 Role of Cryptography 7·3
- 7.1.3 Limitations 7·6

7.2 BASIC CRYPTOGRAPHY 7·6
- 7.2.1 Early Ciphers 7·6
- 7.2.2 More Cryptic Terminology 7·8
- 7.2.3 Basic Cryptanalysis 7·8
- 7.2.4 Brute Force Cryptanalysis 7·9
- 7.2.5 Monoalphabetical Substitution Ciphers 7·11
- 7.2.6 Polyalphabetical Substitution Ciphers 7·12
- 7.2.7 The Vigenère Cipher 7·13
- 7.2.8 Early-Twentieth-Century Cryptanalysis 7·14
- 7.2.9 Adding up XOR 7·15

7.3 DES AND MODERN ENCRYPTION 7·17
- 7.3.1 Real Constraints 7·17
- 7.3.2 One-Time Pad 7·17
- 7.3.3 Transposition, Rotors, Products, and Blocks 7·18
- 7.3.4 Data Encryption Standard 7·20
- 7.3.5 DES Strength 7·20
- 7.3.6 DES Weakness 7·21

7.4 PUBLIC KEY ENCRYPTION 7·21
- 7.4.1 Key-Exchange Problem 7·23
- 7.4.2 Public Key Systems 7·23
- 7.4.3 Authenticity and Trust 7·26
- 7.4.4 Limitations and Combinations 7·27

7.5 PRACTICAL ENCRYPTION 7·27
- 7.5.1 Communications and Storage 7·28
- 7.5.2 Securing the Transport Layer 7·28
- 7.5.3 X.509v3 Certificate Format 7·31

7.6 BEYOND RSA AND DES 7·36
- 7.6.1 Elliptic Curve Cryptography 7·36
- 7.6.2 RSA Patent Expires 7·37
- 7.6.3 DES Superseded 7·38
- 7.6.4 Quantum Cryptography 7·39
- 7.6.5 Snake Oil Factor 7·44

7.7 STEGANOGRAPHY 7·44

7.8 FURTHER READING 7·45

7.9 NOTES 7·46

7.1 INTRODUCTION TO CRYPTOGRAPHY. The ability to transform data so that they are accessible only to authorized persons is just one of the many valuable services performed by the technology commonly referred to as encryption. This technology has appeared in other chapters, but some readers may not be familiar with its principles and origins. The purpose of this chapter is to explain encryption technology in basic terms and to describe its application in areas such as file encryption, message scrambling, authentication, and secure Internet transactions. This is not a theoretical

or scientific treatise on encryption, but a practical guide for those who need to employ encryption in a computer security context.

Organizations around the world increasingly rely on cryptography to communicate securely and to store information safely. Typically, the algorithms used by the Department of Defense (DoD) organizations are employed and maintained for many years. For example, the Data Encryption Standard (DES) has been used in some form for over 20 years.[1]

This chapter is a brief overview of cryptography and its practical applications to the needs of normal business users, as distinct from the needs of high-security government agencies. A thorough examination of the mathematics that are the foundation of these topics is beyond the scope of this chapter, but we provide suggested readings for further study.

7.1.1 Terminology. This list of basic terms will be helpful for readers as they continue through this chapter:

Algorithm—a finite list of well-defined instructions for accomplishing some task that, given an initial state, will terminate in a defined end state.

Cipher—the core algorithm used to encrypt data. A cipher transforms plaintext into ciphertext that is not reversible without a key.

Ciphertext—text in encrypted form, as opposed to the plain text. We show ciphertext in UPPERCASE throughout this chapter.

Codes—a list of equivalences (a *codebook*) allows the substitution of meaningful text for words, phrases, or sentences in an innocuous message; for example, "I will buy flowers for Mama tomorrow for her party at 7 pm" might be decoded to mean "Launch the attack on the mother ship next week on Sunday."

Decrypt/Decipher—the process of retrieving the plaintext from the ciphertext.

Encrypt/Encipher—to alter plaintext using a secret code so as to be unintelligible to unauthorized parties.

Key—a word or system for solving a cipher or code.

Plaintext—the original message to be encoded or enciphered. We show plaintext in lowercase throughout this chapter.

The science of cryptology (sometimes abbreviated as *crypto*) is the study of secure communications, formed from the Greek words $\kappa\rho\upsilon\pi\tau\omicron\sigma$ (*kryptos*), meaning "hidden," and $\lambda\omicron\gamma\omicron\sigma$ (*logos*), "word." More specifically, it is the study of two distinct, yet highly intertwined, fields of study: cryptography and cryptanalysis. Cryptography is "the science of coding and decoding messages so as to keep these messages secure."[2] Cryptanalysis is the art and science of "cracking codes, decoding secrets, violating authentication schemes, and in general, breaking cryptographic protocols,"[3] all without knowing the secret key. Systems for encrypting information are referred to as *cryptosystems*.

Systems for encrypting information may also be referred to as *ciphersystems*, from *cipher*, meaning "zero," or "empty" (a word rooted in the Arabic *sifr*). Terms using cipher and crypto are interchangeable, with some authors preferring cipher to avoid the religious and cultural connotations of *crypt*, a word with the same root as "encryption." Thus, encryption may be referred to as encipherment, decryption referred to as decipherment, and so on.

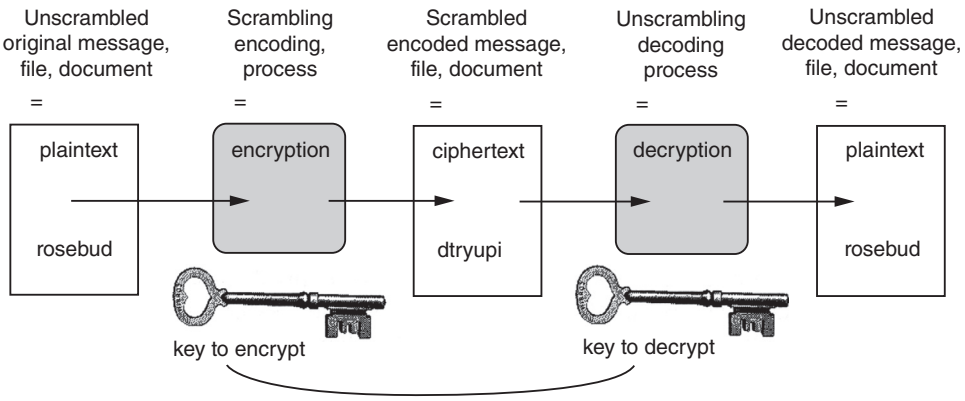


EXHIBIT 7.1 Diagram of Cryptographic Terms

The most obvious use of encryption is to scramble the contents of a file or message, using some form of shared secret as a *key*. Without the key, the scrambled data remain hidden and cannot be unscrambled or *decrypted*. The total number of possible keys for an encryption algorithm is called the *keyspace.* The keyspace is a function of the length of the key and the number of possible values in each position of the key. For a *keylength* of *n* positions, with each position having *v* possible values, then the keyspace for that key would be v^n. For example, with three positions and two values per position (e.g., 0 or 1), the possible keys would be 000, 001, 010, 011, 100, 101, 110, and 111 for a total keyspace of 8.

In cryptographic terms, the contents of a file before encryption are *plaintext,* while the scrambled or encoded file is known as *ciphertext* (see Exhibit 7.1). As a field of intellectual activity, cryptology goes back many millennia. Used in ancient Egyptian, China, and India, it was discussed by the Greeks and regularly employed by the Romans. The first European treatise on the subject appeared in the fourteenth century. The subject assumed immense historic importance during both world wars. The British success in breaking codes that the Germans used to protect military communications in World War II was a major factor in both the outcome of the war and in the development of the first electronic computer systems.

Since then, cryptography and computer science have developed hand in hand. In 1956, the United States National Security Agency (NSA), the U.S. Government department in charge of monitoring the worldwide flow of information, began funding improvements in computer hardware, pumping some $25 million into Project Lightning. This five-year development effort, intended to produce a thousand-fold increase in computing power, resulted in over 150 technical articles. It also gave rise to more than 300 patent applications and succeeded in advancing the frontiers of hardware design. The NSA, based in Fort Meade, Maryland, was also involved in the creation of DES as the commercial encryption standard for much of the last 20 years. Today, the NSA is widely believed to have the world's largest collection of supercomputers and the largest staff of cryptanalysts.

7.1.2 Role of Cryptography. The central role of cryptography in computer security is ensuring the confidentiality of data. But cryptography can support other pillars of computer security, such as integrity and authenticity. This section looks at the different roles of cryptography.

7.1.2.1 Confidentiality. The role of encryption in protecting confidentiality can be seen in a classic definition of encryption: "Encryption is a special computation that operates on messages, converting them into a representation that is meaningless for all parties other than the intended receiver."[4]

Much of the literature on cryptography discusses the technology in terms of ensuring the confidentiality of messages, but this is functionally equivalent to protecting the confidentiality of data. The use of the term "message" reflects the traditional use to which cryptography has been put, both before and after the advent of computers. For example, Julius Caesar encrypted messages to Cicero 2,000 years ago, while today messages between a Web browser and a Web server are encrypted when performing a "secure" transaction.

When applying cryptography to computer security, it is sometimes appropriate to substitute the term "files" for "messages." For example, hard drive encryption programs protect data files stored on a hard drive. However, data files take the form of messages when they are transferred from one computer to another, across a network, the Internet, or via phone lines. Practically speaking, data being transferred in this manner are exposed to a different set of dangers from those that threaten data residing on a computer in an office. Thus, the use of encryption to render files useless to anyone other than an authorized user is relevant both to files in transit and to those that reside on a server or a stand-alone computer, particularly when the latter is a laptop, notebook, or PDA.

7.1.2.2 Integrity. In the second half of the last century, following the advent of programmable computer systems, the ability of cryptography to transform data was applied in many new and interesting ways. As will be seen in a moment, many cryptographic techniques use a lot of mathematical calculation. The ability of computers to perform many calculations in a short period of time greatly expanded the usefulness of cryptography, and also inspired the development of ever-stronger ciphersystems.

Maintaining the integrity of data is often as important as keeping them confidential. When writing checks, people take pains to thwart alteration of the payee or the amount. In some cases, integrity is more important than confidentiality. Changing the contents of a company press release as it passes from the company to the press could have serious consequences. It is not only human actions that threaten data integrity; mechanical failures and logical errors can also change data. It is vital that such changes be detected, as was discussed in Chapter 4 of this *Handbook*, where it was observed that "[a]ll data movements and translations increase the likelihood of internal error, and for this reason parity checks and validity tests have become indispensable."

That chapter covered the role of parity bits for error detection, the function of redundancy checks, and the use of checksums to provide a modification-detection capability. A type of cryptographic hash or checksum called a Message Authentication Code (MAC) can protect against intentional, but unauthorized, data modification as well as against accidental modification. A MAC is calculated by applying a cryptographic algorithm and a secret value, called the *key*, to the data. The data are later verified by applying the cryptographic algorithm and the same secret key to the data to produce another MAC; this MAC then is compared to the initial MAC. If the two MACs are equal, then the data are considered authentic (see diagram in Exhibit 7.2, which uses the public key cryptosystem, discussed later). Otherwise, an unauthorized modification is assumed (any party trying to modify the data without knowing the key would not know how to calculate the appropriate MAC corresponding to the altered data).

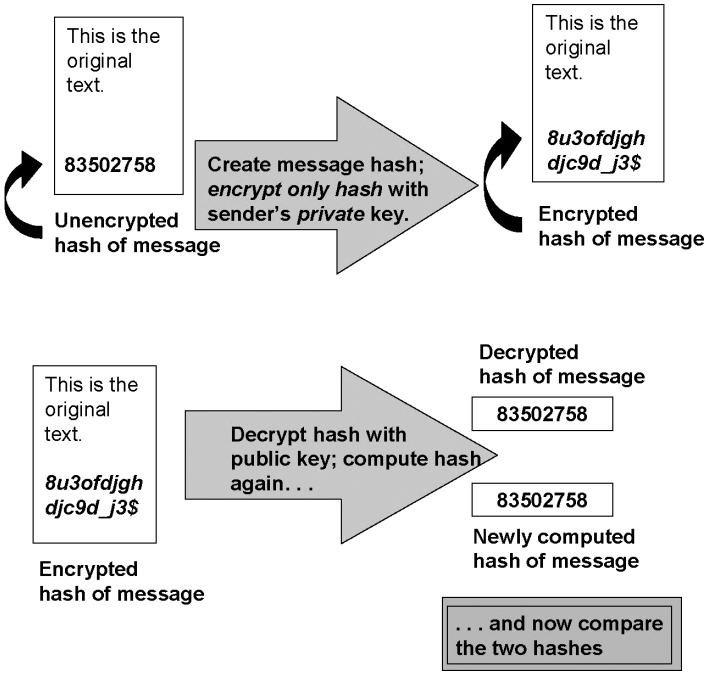


EXHIBIT 7.2 Message Authentication Code Using Public Key Cryptosystem
Source: Copyright © 2008 M. E. Kabay. Used with permission.

7.1.2.3 Authentication. In the context of computer security, authentication is the ability to confirm the identity of users. For example, many computers now ask users to log on before they can access data. By requesting a user name and password, systems attempt to assure themselves that only authentic users can gain access. However, this form of authentication is limited—it merely assures that the person logging on is someone who knows a valid user name and password pair. Cryptography plays a very important role in efforts to ensure stronger authentication, from encrypting the password data to the creation and verification of electronic identifiers such as digital signatures. These will be described in more detail later in this chapter, along with the differences between public key and private key cryptography, both of which may be used in these schemes.

Using a public key system, documents in a computer system can be electronically signed by applying the originator's private key to the document. The resulting digital signature and document then can be stored or transmitted. The signature can be verified using the public key of the originator. If the signature verifies properly, the receiver has confidence that the document was signed using the private key of the originator and that the message had not been altered after it was signed. Because private keys are known only to their owner, it is also possible to verify the originator of the information to a third party.

7.1.2.4 Nonrepudiation. An aspect of computer security that has increased greatly in significance, due to the growth in internetwork transactions, is nonrepudiation. For example, if someone places an electronic order to sell stocks that later increase

in value, it is important to prove that the order definitely originated with the individual who placed it. Made possible by public key cryptography, nonrepudiation helps ensure that the parties to a communication cannot deny having participated in all or part of the communication.

7.1.3 Limitations. One role that cryptography cannot fill is defense against data destruction. Although encryption does not assure availability, it does represent a very valuable extra line of defense for computer information when added to physical security, system access controls, and secure channels of communication. Indeed, when computers are mobile, or data are being communicated over insecure channels, encryption may be the main line of defense. However, even though applied cryptography can provide computer users with levels of security that cannot be overcome without specialized knowledge and powerful computers, encryption of data should not be thought of as an alternative to, or substitute for, system access control. According to Seberry and Pieprzyk[5], the role of cryptography is to protect "information to which illegal access is possible and where other protective measures are inefficient."

Encryption-based file access controls should be a third barrier after site and system access controls, if for no other reason than that encryption systems alone do little to prevent people deleting files.

7.2 BASIC CRYPTOGRAPHY. The aim of cryptography is to develop systems that can encrypt plaintext into ciphertext that is indistinguishable from a purely random collection of data. This implies that all of the possible decrypted versions of the data except one will be hopelessly ambiguous, with none more likely to be correct than any of the others. One of the simplest ways to create ciphertext is to represent each character or word in the plaintext by a different character or word in the ciphertext, such that there is no immediately apparent relationship between the two versions of the same text.

7.2.1 Early Ciphers. It is believed that the earliest text to exhibit the baseline attribute of cryptography, having a slight modification of the text, occurred in Egypt nearly 4,000 years ago. A scribe used a number of unusual symbols to confuse or obscure the meaning of the hieroglyphic inscriptions on the tomb of a nobleman named Khnumhotep II.[6]

It is also believed that the first effective military use of cryptography was a simple transposition cipher (see Section 7.3.3) by the Spartans, who "as early as 400 BCE employed a cipher device called the scytale for secret communication between military commanders."[7] The scytale was a cylindrical or tapered stick with a thin strip of leather or parchment wrapped around it spirally.[8] The message to be hidden was written lengthwise with no blank spaces. When unraveled, the parchment appeared to hold nothing but random letters. To read the parchment, the recipient had to have a stick with exactly the same dimensions as the sender. The distribution of appropriate decoding scytales took place before the military commanders departed for the field.[9] For example, a particular combination of stick and strip could allow the cleartext (shown in lowercase):

```
atheniantroopswithinonedaysmarchofromebereadynow
```

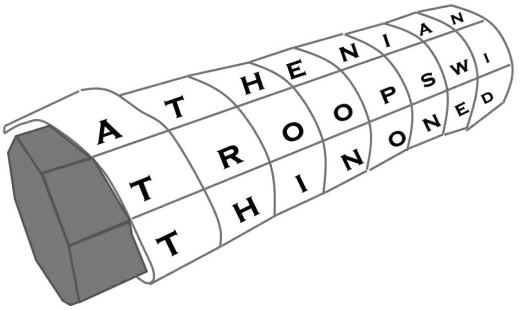


EXHIBIT 7.3 Scytale in Use
Source: Copyright © 2008 M. E. Kabay. Used with permission.

to be broken into up to six rows of eight letters to be written across the rolled-up strip, in this way:

athenian
troopswi
thinoned
aysmarch
ofromebe
readynow

The message might appear on the scytale as shown schematically in Exhibit 7.3.

Reading the unwrapped strip without the stick would produce this ciphertext (shown in uppercase):

```
ATTAORTRHYFEHOISREEONMODNPOAMYISNRENAWECBONIDHEW
```

"The first attested use of [a substitution cipher] in military affairs comes from the Romans."[10] During that time, Julius Caesar encoded all his messages by simply replacing every letter with the letter three positions away. For example, the letter *a* would become the letter *d*, the letter *b* would become the letter *e*, and so on. Now called the Caesar cipher, this scheme is best-known of all the monoalphabetic algorithms (see Section 7.2.5).[11] Consider the Caesar cipher illustrated in the next comparison using the modern English alphabet, with the letters of the alphabet simply shifted three places.

```
Plaintext:  abcdefghijklmnopqrstuvwxyz
Ciphertext: DEFGHIJKLMNOPQRSTUVWXYZABC
```

To encrypt a message, the sender finds each letter of the message in the plaintext alphabet and uses the letter below it in the ciphertext alphabet. Thus, the clear message:

```
Plaintext: beware the ides of march
```

is transformed into the encrypted message:

```
Ciphertext: EHZDUH WKH LGHV RI PDUFK
```

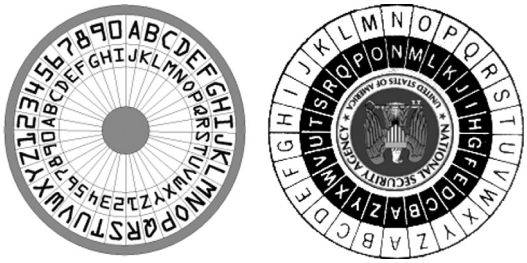


EXHIBIT 7.4 Code Wheels and the NSA Seal

This type of cipher is known as a *substitution cipher.* Although the Caesar cipher is relatively simple, substitution ciphers can be very powerful. Most examples of the Caesar cipher shift the alphabet three places, as shown, so that the ciphertext line begins with *d*, but some authors suggest Caesar might have used other numbers, so the term "Caesar cipher" is used for all ciphers that conform to this algorithm (an *algorithm* being a formula or recipe for solving a problem).

This level of encryption might seem rudimentary, but it is an important starting point for much that follows. For example, one way to visualize the Caesar cipher is as a pair of rings, one inside the other, as shown in Exhibit 7.4. Both circles contain the letters of the alphabet. If one is rotated relative to the other, the result is a cipher wheel, something well suited to automation. Eventually this happened, at first mechanically, then electrically, and today digitally. Automation facilitates repetition, and messages encrypted with a substitution cipher can be more difficult to decipher if multiple different substitutions are used. Thus, the code wheel earned a place in the seal of the NSA, the U.S. government agency most influential in the development of encryption.

7.2.2 More Cryptic Terminology. The key or password for the Caesar cipher presented in the last section is the number of places the alphabet has been shifted, in this case three. Because this key must be kept private in order for the message to remain protected, it must be delivered to the recipient for the message to be decoded, or decrypted, back to plaintext. That is why the Caesar cipher is described as a *private key* algorithm and also a *symmetrical encryption* algorithm, the same private key being used to encrypt and decrypt the message. Algorithms of this type can be defeated by someone who has the key, an encrypted message, and knowledge of the algorithm used. This might sound like a statement of the obvious; however, as will be seen later in this chapter, there are encryption algorithms that use keys that can be openly exchanged without rendering the encrypted data accessible. Knowledge of the algorithm used can often be derived, or reverse-engineered, by analysis of its output.

Another seemingly obvious fact is that when a private key cipher is used in an effort to achieve confidentiality, one problem is swapped for another. The problem of exchanging messages while keeping the contents from unintended recipients is replaced by the problem of exchanging keys between sender and receiver without disclosing the keys. This new problem is known as the *key-exchange problem*. The key-exchange problem will be examined in more detail later.

7.2.3 Basic Cryptanalysis. "The first people to understand clearly the principles of cryptography and to elucidate the beginnings of cryptanalysis were the Arabs."[12] By the fifteenth century, they had discovered the technique of letter frequency distribution analysis and had successfully decrypted a Greek message on its way to the

Byzantine Emperor.[13] In 1492, a man known as al-Kalka-shandi described this technique in an encyclopedia. He also described several cryptographic techniques, including substitution and transposition ciphers.[14]

Returning to the Caesar cipher, consider how this code could be broken using the science of cryptanalysis. When examined for a length of time, this particular code is fairly transparent. As soon as several letters are identified correctly, the rest fall into place. For example, because "the" is the most common three-letter word in the English language, testing "XLI" against "the" reveals that each letter of plaintext has a fixed relationship to the ciphertext: a shift of three to the right.

If that difference is applied to the rest of the message, the result is a piece of plaintext that is intelligible and thus assumed to be the correct solution to the problem. However, even in this simple example several sophisticated processes and assumptions are at work; they deserve closer attention before looking at more complex codes. First, the test of "the" against "XLI" assumes that the plaintext is English and that the attacker has some detailed knowledge of that language, such as the frequency of certain words. Second, it is assumed that the ciphertext follows the plaintext in terms of word breaks. Typically, this is not the case. Ciphertext usually is written in blocks of letters of equal length to further disguise it, as in:

```
Ciphertext: EHZDU HWKHL GHVRI PDUFK
```

When the recipient of the message decrypts it, the result, while not exactly easy reading, is nevertheless entirely intelligible:

```
Plaintext: bewar ethei desof march
```

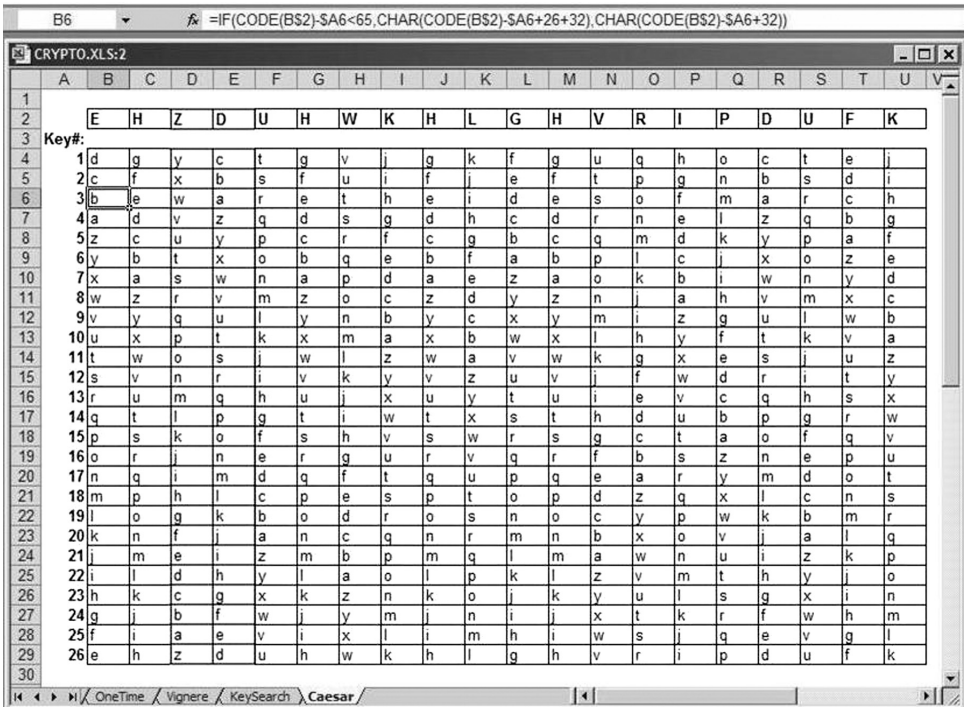
Also note the convention of ignoring the case of individual letters and placing all plaintext in lowercase while all ciphertext is in capitals.

7.2.4 Brute Force Cryptanalysis. The next thing to note about the Caesar cipher is that, using the English alphabet, there are 26 possible keys. This means that someone intercepting the encrypted message could mount a standard form of attack known as *brute force cryptanalysis*. This method runs possible keys through the decryption algorithm until a solution is discovered. Statistically speaking, the correct key is reached after testing only half of all possible keys. In Exhibit 7.5, a spreadsheet table details a brute force attack on the Caesar ciphertext. In the example, the plaintext appears in line 6, Key #3.

Note that three items of information are required for this attack, and all three of them are relevant to encryption on personal computers:

1. A knowledge of the encryption algorithm used
2. The number of possible keys
3. The language of the plaintext

Using a computer in an office is somewhat different from sending messages on the field of battle (at least on a good day). Unlike an enemy spy, someone who is attempting to gain unauthorized access to data already has a fairly good idea of which algorithm is being used. (There are relatively few in use, and they often are directly associated with particular applications). This takes care of the first item. The primary obstacle to a brute force attack is the second item, number of keys. In the case of the Caesar



B6 =IF(CODE(B$2)-$A6<65,CHAR(CODE(B$2)-$A6+26+32),CHAR(CODE(B$2)-$A6+32))

CRYPTO.XLS:2

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
1																					
2		E	H	Z	D	U	H	W	K	H	L	G	H	V	R	I	P	D	U	F	K
3	Key#:																				
4	1	d	g	y	c	t	g	v	j	g	k	f	g	u	q	h	o	c	t	e	j
5	2	c	f	x	b	s	f	u	i	f	j	e	f	t	p	g	n	b	s	d	i
6	3	b	e	w	a	r	e	t	h	e	i	d	e	s	o	f	m	a	r	c	h
7	4	a	d	v	z	q	d	s	g	d	h	c	d	r	n	e	l	z	q	b	g
8	5	z	c	u	y	p	c	r	f	c	g	b	c	q	m	d	k	y	p	a	f
9	6	y	b	t	x	o	b	q	e	b	f	a	b	p	l	c	j	x	o	z	e
10	7	x	a	s	w	n	a	p	d	a	e	z	a	o	k	b	i	w	n	y	d
11	8	w	z	r	v	m	z	o	c	z	d	y	z	n	j	a	h	v	m	x	c
12	9	v	y	q	u	l	y	n	b	y	c	x	y	m	i	z	g	u	l	w	b
13	10	u	x	p	t	k	x	m	a	x	b	w	x	l	h	y	f	t	k	v	a
14	11	t	w	o	s	j	w	l	z	w	a	v	w	k	g	x	e	s	j	u	z
15	12	s	v	n	r	i	v	k	y	v	z	u	v	j	f	w	d	r	i	t	y
16	13	r	u	m	q	h	u	j	x	u	y	t	u	i	e	v	c	q	h	s	x
17	14	q	t	l	p	g	t	i	w	t	x	s	t	h	d	u	b	p	g	r	w
18	15	p	s	k	o	f	s	h	v	s	w	r	s	g	c	t	a	o	f	q	v
19	16	o	r	j	n	e	r	g	u	r	v	q	r	f	b	s	z	n	e	p	u
20	17	n	q	i	m	d	q	f	t	q	u	p	q	e	a	r	y	m	d	o	t
21	18	m	p	h	l	c	p	e	s	p	t	o	p	d	z	q	x	l	c	n	s
22	19	l	o	g	k	b	o	d	r	o	s	n	o	c	y	p	w	k	b	m	r
23	20	k	n	f	j	a	n	c	q	n	r	m	n	b	x	o	v	j	a	l	q
24	21	j	m	e	i	z	m	b	p	m	q	l	m	a	w	n	u	i	z	k	p
25	22	i	l	d	h	y	l	a	o	l	p	k	l	z	v	m	t	h	y	j	o
26	23	h	k	c	g	x	k	z	n	k	o	j	k	y	u	l	s	g	x	i	n
27	24	g	j	b	f	w	j	y	m	j	n	i	j	x	t	k	r	f	w	h	m
28	25	f	i	a	e	v	i	x	l	i	m	h	i	w	s	j	q	e	v	g	l
29	26	e	h	z	d	u	h	w	k	h	l	g	h	v	r	i	p	d	u	f	k
30																					

OneTime | Vignere | KeySearch | Caesar

EXHIBIT 7.5 Brute Force Attack on the Caesar Cipher

cipher, the number of possible keys is relatively small, so the work involved in carrying out the attack can be completed very quickly, which is highly significant. Time is often the most important factor in practical cryptanalysis. Being able to decrypt messages within 24 hours is of little use if the information pertains to events that are measured in minutes, such as orders to buy and sell stock, or to launch air raids. If the cipher consisted entirely of random letter substitutions, like this:

```
Plaintext:  abcdefghijklmnopqrstuvwxyz
Ciphertext: UTWFRAQOYSEDCKJVBXGZIPHLNM
```

The number of possible keys (the *keyspace*) is now 26!, or $\sim 4.03 \times 10^{26}$, which looks even more daunting when written out:

```
403,291,461,126,606,000,000,000,000
```

Imagine a brute force attack using a computer that can perform 1 million decryptions per microsecond (considerably more number crunching than the average personal computer can perform). Using a single processor, it could take over 10 million years to execute a brute force attack on this code. Fortunately for the code breaker, there are other ways of cracking substitution ciphers, as discussed in a moment. The point is that, while brute force attacks are possible, they are not always practical.

Although it is true that by the central limit theorem of statistics, the most likely number of trials required to hit on the correct key is one-half the total keyspace, the average reduction by a factor of 2 is negligible in the face of computational periods measured in years and the difficulty of identifying cleartext in the morass of incorrect decryptions.

Functionally, brute force attacks depend on knowing which encryption algorithm is behind the ciphertext. Practically, they depend on the feasibility of successes within an appropriate time frame. They also depend on the third item of information in the list above: knowledge of the language of the plaintext. The solution to the Caesar cipher in Exhibit 7.5 tends to jump out because it is closer to plain English than any of the other solutions. However, without knowing what constitutes plaintext, a brute force attack will, at best, be inefficient, and, at worst, unsuccessful. This part of cryptanalysis, recognizing a positive result, is less amenable to automation than any other. The difficulty is compounded by encryption of purely numerical results where the correct cleartext can be impossible to determine without extensive additional knowledge.

7.2.5 Monoalphabetical Substitution Ciphers. Both the Caesar cipher and the random substitution cipher shown are examples of monoalphabetic ciphers. This means that one letter of ciphertext stands for one letter of plaintext. This renders such codes susceptible to an attack quite different from brute force. Suppose a customs officer attempts to discover when and how an illegal weapons shipment will be entering the country. The following message is intercepted:

```
YZYGJ KZORZ OYXZR RKZRK XUXRJ XRZXU YKQQQ
```

The person who encoded this text clearly substituted new letters for the original letters of the message. To the experienced code breaker or cryptanalyst, the task of deciphering this message is quite a simple one. First count how many times each letter occurs in the text. This produces a list like this:

```
Ciphertext: R Z X Y K J U O G
Frequency:  6 6 5 4 4 2 2 2 1
```

Note that the last three letters are discounted as they are merely filling out the five-letter grouping. Next refer to a table of frequencies, which shows the relative frequency with which the letters of the alphabet occur in a specific language or dialect of that language. One such list is shown in Exhibit 7.6. This list was created for this example and proposes that the most commonly used letters in English in descending order of frequency are *e, t, r,* and so on. The actual order is more likely to be *e, t, a, i, o, n, s, h, r, d, l, u*, the order of keys on the English Linotype machine from the nineteenth century, although the precise order of frequencies can vary according to the region of origin or subject matter of the text.

Assuming that the original message is in English, a list that matches code letters to plaintext letters is easily derived.

```
Ciphertext: R Z X Y K J U O G
Frequency:  6 6 5 4 4 2 2 2 1
Plaintext:  e t r i n o a h s
```

The result is:

```
Ciphertext: YZYGJ KZORZ OYXZR RKZRK XUXRJ XRZXU YKQQQ
Plaintext:  itiso nthet hirte enten rareo retra inqqq
```

This is readable as "it is on the thirteen ten rare ore train." Although this example obviously was contrived to make a point, it clearly illustrates an important cryptographic

EXHIBIT 7.6 Frequency Lists for English

English by Letter				English by Frequency			
A	7.25	N	7.75	E	12.75	U	3.00
B	1.25	O	7.50	T	9.25	M	2.75
C	3.50	P	2.75	R	8.50	P	2.75
D	4.25	Q	0.50	I	7.75	Y	2.25
E	12.75	R	8.50	N	7.75	G	2.00
F	3.00	S	6.00	O	7.50	V	1.50
G	2.00	T	9.25	A	7.25	W	1.50
H	3.50	U	3.00	S	6.00	B	1.25
I	7.75	V	1.50	D	4.25	K	0.50
J	0.25	W	1.50	L	3.75	Q	0.50
K	0.50	X	0.50	C	3.50	X	0.50
L	3.75	Y	2.25	H	3.50	J	0.25
M	2.75	Z	0.25	F	3.00	Z	0.25

tool that can quickly decipher something that at first seems to be very forbidding. The encryption in the previous example could have been based on a simple substitution cipher. For example, after using the password "TRICK" followed by the regular alphabet minus the letters in the password for the plaintext, the ciphertext is the alphabet written backward:

```
Plaintext:  TRICKABDEFGHJLMNOPQSUVWXYZ
Ciphertext: ZYXWVUTSRQPONMLKJIHGFEDCBA
```

Frequency analysis also works if the substitution is entirely random, as in the example shown earlier, the key for which is entirely random. The specialized tools, such as frequency tables, that are required to break codes point out a basic trade-off: If a basic level of protection is needed, it is easy to get but also easy to break, at least for an expert. The qualification "for an expert" is important because users of encryption need to keep its role in perspective. The salient questions are: Who can gain from decrypting the data, and what means do they have at their disposal? There is no point investing in powerful encryption hardware or software if those likely to attempt to read your files are not particularly sophisticated, dedicated, or well equipped. For example, a person who mails a postcard knows it can be read by anyone who sees it. Envelopes can be used to prevent this, hardly the ultimate in confidentiality, but widely used and relatively successful nonetheless.

7.2.6 Polyalphabetical Substitution Ciphers. Even when the plaintext uses a wider range of letters than the contrived example, substitution ciphers can be cracked by frequency analysis. A powerful technique is to concentrate on the frequency of two-letter combinations, which are known as *digraphs,* the most common of which in English is "TH." One way to counter frequency analysis is to use multiple substitutes for the more frequent letters. This cannot be done with a straightforward alphabetic coding. However, if using numbers for letters, it is possible to assign multiple numbers to some letters, such as 13 17 19 23 for E, which would help dilute the natural frequency of this letter. It would appear that supplying multiple substitutions, known as *homophones,* in proportion to the frequency of each letter would effectively

counter frequency analysis. However, some of the underlying structure of the plaintext still survives, notably digraphs, which the cryptanalyst can use to crack the code.

In Europe during the Middle Ages, advances in cryptography were being made by the Papal States and Italian city-states to protect diplomatic messages. Then, in 1379, an Italian man named Gabriele de Lavinde created the first European manual on cryptography. "This manual, now in the Vatican archives, contains a set of keys for 24 correspondents and embraces symbols for letters, nulls, and several two-character code equivalents for words and names."[15] The nomenclature described by Lavinde's manual "was to hold sway over all Europe and America for the next 450 years."[16]

Several other notable advances emerged in Europe during the period of Lavinde's manual. First, in 1470, Leon Battista Alberti published the first description of a cipher disk.[17] Next, in 1563, Giambattista della Porta provided the first example of a digraphic cipher in which two letters are represented by one symbol.[18]

One method of decreasing the extent to which the structure of the plaintext is reflected in the ciphertext is to encrypt multiple letters of the plaintext. For example, "AR" might be encrypted as "CM." This is the theory behind what is known as the Playfair cipher, which was invented in 1854 by a British scientist, Sir Charles Wheatstone, but that was named after his friend Baron Playfair who fought for its adoption by the British Foreign Office.[19] Although the Playfair cipher remained in use through both world wars, it does not do enough to disguise the plaintext and cannot withstand a concerted frequency analysis.

7.2.7 The Vigenère Cipher. A particularly important technique in the evolution of polyalphabetic ciphers has its roots in the sixteenth century. In 1586, Blaise de Vigenère published a square encryption/decryption table, named after him as the Vigenère Square, and descriptions of the first plaintext and ciphertext autokey systems.[20] The Vigenère cipher involves a table of letters, like the one shown in Exhibit 7.7, that are used with a key to provide different monoalphabetic substitutions as the encryption proceeds through the plaintext. Thus, each letter of the ciphertext has a different relationship with the plaintext, like this:

```
Key:        doomsdaydoomsdaydoomsdaydoomsday
plaintext:  sellentireportfolionowandbuygold
ciphertext: VSZXWQTGJIAZVGYWCMDBFPFBOJIKUKLQ
```

The message is enciphered by looking at the row in the table that begins with the first letter of the key. Then go along that row until the column headed by the first letter of the plaintext. The ciphertext substitution is the letter at that intersection in the table. Thus, row d, column s, yields V. Then proceed to the second letter, and so on. Note that the first time the letter *e* is encrypted the cipher is S, but the second time it is W. The two *l*s in sell are encoded as Z and X, respectively, and so on.

Does this cipher completely obscure the structure of the plaintext? Stallings notes: "If two identical sequences of plaintext letters occur at a distance that is an integer multiple of the keyword length, they will generate identical ciphertext sequences."[21] This means that the cryptanalyst can determine the length of the keyword. Once this is done, the cipher can be treated as a number of monoalphabetic substitutions, that number being equal to the key length. Frequency tables are again brought into play, and the code can be cracked. The cryptographer's response to this weakness is to use a longer key so that it repeats less often. In fact, one technique, *autokey,* invented by

	a	b	c	d	e	f	G	h	i	j	k	l	m	n	o	p	q	r	s	t	u	v	w	x	y	Z
a	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
b	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A
c	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B
d	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C
e	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D
f	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E
g	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F
h	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G
i	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H
j	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I
k	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J
l	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K
m	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L
n	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M
o	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N
p	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
q	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
r	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
s	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
t	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
u	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T
v	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
w	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V
x	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W
y	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	Q	X
z	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y

EXHIBIT 7.7 Vigenère Table

Vigenère, is to form the key from the plaintext itself, together with one code word, like this

```
Key:         doomsdaysellentireportfolionowan
plaintext:   sellentireportfolionowandbuygold
ciphertext: VSZXWQTGJIAZVGYWCMDBFPFBOJILUKLQ
```

This system is very powerful, but it still can be attacked by statistical analysis based on frequencies, because the letters of the plaintext and key share roughly the same frequency distribution. The next level of defense is to use a keyword that is as long as the plaintext but bears no statistical relationship to it. This approach, which is of great cryptographic significance, was not hit upon until the twentieth century arrived, bringing with it binary code and global warfare.

7.2.8 Early-Twentieth-Century Cryptanalysis. The advent of modern cryptography began with the invention and development of the electromagnetic telegraph system and the introduction of the Morse code. Samuel Morse brought a system of dots and dashes that allowed near real–time long-distance communication. He envisioned this system as a means of secure communications. It would be up to others to devise systems to encrypt telegraphic communications. Anson Stager, the supervisor

of the U.S. Military Telegraph during the Civil War, devised 10 ciphers for the Union Army that were never broken by the Confederacy.[22]

The use of telegraphic ciphers and codes continued into the two world wars. In fact, one of the most famous early successes of cryptanalysis prompted the entrance of the United States into World War I. When the war first started, the German transatlantic telegraph cable had been cut by the British, forcing all of Germany's international communications to route through the United Kingdom before being sent on to the Swedish or American transatlantic lines.[23] In 1917, "British cryptographers deciphered a telegram from German Foreign Minister Arthur Zimmermann to the German Minister to Mexico, von Eckhardt."[24] It promised Mexico ownership over territory that belonged to the United States (e.g., California), if Mexico joined the German cause and attacked the United States. The British informed President Wilson of their discovery, giving him a complete copy of the telegram, thus resulting in the United States declaring war on Germany.[25] That telegram has become famous in the history of cryptanalysis as the Zimmermann Telegram.

World War II saw several Allied victories over the Axis powers by use of advanced cryptographic systems. Few of these victories are more widely known and celebrated than the cracking of the German Enigma cipher machine, described next:

> Following the decryption of the Zimmerman Telegram during World War I and the effects that weak ciphers had on that war's outcome, Germany was looking for an unbreakable cipher and was interested in leveraging automation and the use of machinery to replace traditional paper and pencil techniques. The Enigma machine consisted of a basic keyboard, a display that would reveal the cipher text letter, and a scrambling mechanism such that each plain text letter entered as input via the keyboard was transcribed to its corresponding cipher text letter. The machine was modular in design and multiple scrambling disks were employed to thwart attempts at frequency analysis.[26]

A British cryptanalysis group, with the help of a group of Polish cryptanalysts, first broke the Enigma early in World War II, and some of the first uses of computers were for decoding Enigma ciphers intercepted from the Germans. Breaking Enigma was a major victory for the Allies, and in order to keep exploiting it, they kept the fact that they had cracked it a secret.[26]

Thus far, the encryption schemes or devices described have encrypted messages consisting of words and nothing more. However, the emergence of the computer, even in its initial rudimentary form, revolutionized cryptology "to an extent even greater than the telegraph or radio."[27] Most cryptologic advances since World War II have involved, or made use of, computers. In the last few decades, cryptographic algorithms have advanced to the point where computing them by hand would be unfeasible, and only computers can do the required mathematics.[28] Relying on computers has broadened the kind of information that can benefit from encryption. Computers use a unique language that transforms all information stored into bits, each a 1 or a 0.[29] "This, in effect, means that plaintext is binary in form, and can therefore be anything; a picture, a voice, an email or even a video—it makes no difference, a string of binary bits can represent any of these."[30]

7.2.9 Adding up XOR. In 1917, an engineer at AT&T, Gilbert Vernam, was working on a project to protect telegraph transmissions from the enemy. At that time, teletypewriters were used, based on a version of Morse code called Baudot code, after its French inventor. In Baudot code, each character of the alphabet is allotted five units, each of which is either an electrical current or absence of current, known as a

mark or a space. For example, the letter *a* is represented by *mark, mark, space, space, space.* In binary terms, each unit constitutes a bit that is either 0 or 1 (the five-bit code for *a* would be 11000). This system of pulses allowed teletype machines to convert text to and from telegraph signals using a keyboard and punched paper tape for input (a hole represents a mark because it allows the reading device to make electrical contact and create a pulse, whereas a space is represented by leaving the paper intact). Anyone with a suitable machine could intercept and read the transmission.

The 32 possible combinations (2^5) in this code were assigned to the 26 letters plus six "shunts" that did various things like shift to capitals or go down to the next line. Vernam's brilliant idea was to use a tape of random characters in Baudot code as a key that could be electromechanically added to the plaintext. Kahn describes the method of addition like this:

> If the key and plaintext pulses are both marks or both spaces, the ciphertext pulse will be a space. If the key pulse is a space and the plaintext pulse is a mark, or vice-versa (in other words, if the two are different), the ciphertext will be a mark.[31]

Today, this is known as Exclusive-Or, sometimes referred to as bit-wise XOR or just XOR for short (see Exhibit 7.8). XOR is widely used in computerized encryption schemes. Consider what happens when encoding the letter *a* using *B* as the key:

```
Plaintext:  1 1 0 0 0 (=a)
Key:        1 0 0 1 1 (=B)
Ciphertext: 0 1 0 1 1
```

In the first column, 1 + 1 = 0, as indicated in Exhibit 7.8. To decipher the encrypted character, simply perform the same operation, but add the ciphertext to the key:

```
Ciphertext: 0 1 0 1 1
Key:        1 0 0 1 1 (=B)
Plaintext:  1 1 0 0 0 (=a)
```

At the time of its discovery, the significance of this method lay in its capacity for automation. The operator could feed the plaintext and key tapes into the teletype machine, and it would transmit an encrypted message with no further human input. No offline preparation was required. Furthermore, as long as the receiver had the key tape, the teletype at the receiving end automatically printed out plaintext. This made Vernam's system the first to integrate encryption into the communication process, an essential feature of encryption systems for today's computer-based communications.

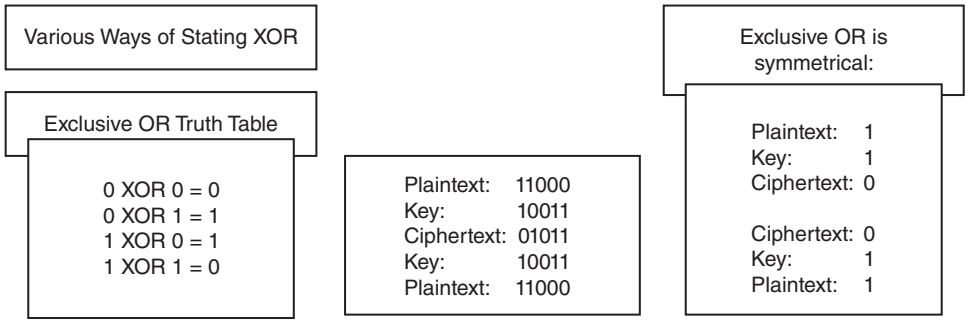


Exhibit 7.8 Diagram of XOR

7.3 DES AND MODERN ENCRYPTION. Although the use of XOR predated computers, the fact that it worked so well with binary code ensured that it would become an essential item in the modern cryptographer's toolkit. And so the focus of this chapter turns to modern cryptography and two of the most widely used cryptosystems today. The first is Data Encryption Standard (DES) and the second is Rivest, Shamir, Adleman (RSA).

7.3.1 Real Constraints. As the preceding overview of the evolution of encryption suggests, major advances, which are few and far between, often are linked with the individuals who made them, such as Vigenère, Playfair, and Vernam, none of whom had the benefit of computers. Today's computerized encryption schemes typically employ a number of classic techniques that, when combined, eliminate or minimize the shortcomings of any single method. Several techniques will be discussed here, including transposition and rotors, that point the way to the most widely used encryption scheme to date: DES. First, however, consider the practical problems encountered by Vernam's otherwise brilliant scheme.

Vernam proposed a key that was a long series of random characters. This was coded on a loop of paper tape that eventually repeated (the tape held about 125 characters per foot). The length of the key made cryptanalysis of intercepted messages extremely difficult, but not impossible, because eventually the key repeated. With sufficient volume of ciphertext, the code would yield to frequency analysis. (Bear in mind that during time of war, or even military exercises, hundreds of thousands of words may be encrypted per day, providing a solid basis for cryptanalysis.)

7.3.2 One-Time Pad. Several improvements then were suggested to avoid the impracticality of simply creating longer and longer key tapes. Another AT&T engineer, Lyman Morehouse, suggested using two key tapes of about eight feet in length, containing some 1,000 characters, to generate over 999,000 combinations of characters that could be fed into the encryption process as the key. This was an improvement in terms of practicality and security, but, as Major Joseph Mauborgne of the U.S. Army Signal Corps pointed out, heavy message traffic encrypted in this way still could be decoded. It was Mauborgne who realized that the only unbreakable cipher would use keys that are, as Kahn puts it "endless and senseless."[32] Thus he came up with what we know as the one-time system, the one unbreakable encryption scheme.

The one-time system sometimes is referred to as a *one-time pad*,[33] because this is the way it has been deployed by intelligence agents in the field. The agent is issued a pad that aligns columns and rows of entirely random characters, as shown in Exhibit 7.9. The first letter of the plaintext is encrypted using the appropriate ciphertext from row 1, the second letter is encrypted from row 2, and so on. The result is ciphertext that contains no statistical relationship to the plaintext. When the message is encrypted the pad is destroyed. The recipient, who has a copy of the pad, uses it to reverse the process and decrypt the message.

The one-time pad essentially is a polyalphabetic substitution cipher, but with the same number of alphabets as there are characters in the message, thus defeating any kind of frequency analysis. A brute force attack is defeated by the fact that every possible result is as statistically significant as every other. As Kahn points out, a four-letter group of ciphertext could just as easily yield *kiss, fast, slow,* or any other possible four-letter combination.

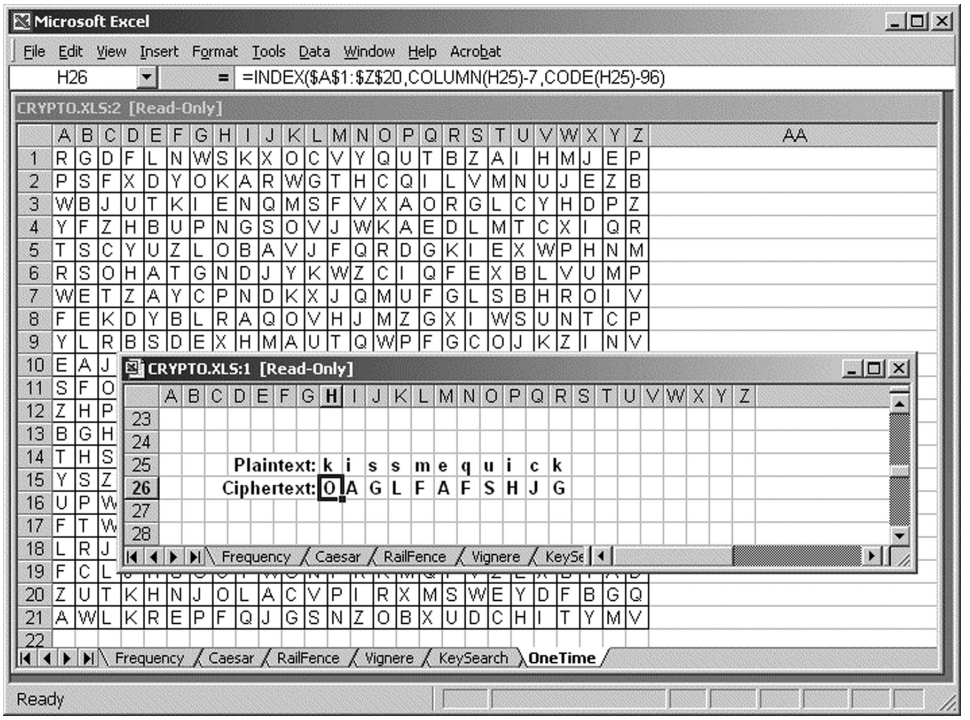


EXHIBIT 7.9 One-Time Pad

So why is the unbreakable one-time system not in universal use? Well, it remains a favorite of intelligence agents in the field who have an occasional need to send short messages. However, for large-scale commercial or military encryption, it fails to solve the key size problem that Vernam's system brought to light. The key has to be as large as the total volume of encrypted information, and there is a constant demand for new keys. Furthermore, both sender and receiver have to hold and defend identical copies of this enormous key.

7.3.3 Transposition, Rotors, Products, and Blocks. A completely different technique from substitution is *transposition.* Instead of substituting ciphertext characters for plaintext, the transposition cipher rearranges the plaintext characters. The simplest example is referred to as *rail fence*. For example, to encrypt "sell entire portfolio now and buy gold" each character is written on alternate lines, like this:

```
sletrprflooadugl
elnieotoinwnbyod
```

which results in this ciphertext:

```
SLETRPRFLOOADUGLELNIEOTOINWNBYOD
```

So far, this does not present a serious challenge. More challenging is the next transposition into rows and columns that are numbered by a key (in this case, 37581426)

so that the first set of ciphertext characters are under 1, the second under 2, and so on:

Key:	3	7	5	8	1	4	2	6
Plaintext:	s	e	l	l	e	n	t	i
	r	e	p	o	r	t	f	o
	l	i	o	n	o	w	a	n
	d	b	u	y	g	o	l	d

```
Ciphertext: EROGTFALSRLDNTWOLPOUIONDEEIBLONY
```

Although more complex, this transposition will still yield to cryptanalysis because it retains the letter frequency characteristics of the plaintext. The analyst also would look for digraphs and trigraphs while playing around with columns and rows of different length. (Kahn describes French code breakers during World War I literally cutting text into strips and sliding them up and down against each other to break German transposition ciphers.)

What makes transposition difficult to decipher is additional stages of encryption. For example, if the previous ciphertext is run through the system again, using the same key, all semblance of pattern seems to disappear.

Key:	3	7	5	8	1	4	2	6
Plaintext:	e	r	o	g	t	f	a	l
	s	r	l	d	n	t	w	o
	l	p	o	u	i	o	n	d
	e	e	i	b	l	o	n	y

```
Ciphertext: TNILAWNNESLEFTOOOLOILODYRRPEGDUB
```

The development of increasingly complex multiple-transposition ciphers pointed out the positive effects of multiple stages of encryption, which also apply to substitution ciphers. The prime examples of this are the rotor machines used by the Germans and Japanese in World War II. Some of the insights gained during the attack on German codes, such as Alan Turing's 1940 work on the application of information statistics to cryptanalysis, were considered so important that they remained classified for more than 50 years.

Although they eventually were defeated by Allied cryptanalysts, electromechanical systems such as Enigma were not only the most sophisticated precomputer encryption systems, but the effort to crack them was also a major catalyst in the development of computer systems themselves. When people started applying computer systems to code making rather than code breaking, they quickly hit on the idea of chopping plaintext into pieces, or blocks, for easier handling. The term "block cipher" is used to describe ciphers that encrypt one block (e.g., 8 bytes of data) at a time, one block after another. Another result of computerizing the encryption process is a class of ciphers known as *product ciphers.* A product cipher has been defined as "a block cipher that iterates several weak operations such as substitution, transposition, modular addition/multiplication [such as XOR], and linear transformation."[34]

The mathematics of product ciphers are beyond the scope of this chapter, but it is useful to note that "[n]obody knows how to prove mathematically that a product cipher is completely secure . . . [A] product cipher should act as a 'mixing' function

which combines the plaintext, key, and ciphertext in a complex nonlinear fashion."[35] The parts of the product cipher that perform the rounds of substitution are referred to as *S-boxes.* The product cipher called Lucifer has two of these S-boxes, while DES encryption has eight S-boxes. The ability of a product cipher to produce truly random, nonlinear ciphertext depends on careful design of these S-boxes.

Examples of modern product ciphers include Lucifer (developed by IBM), DES (developed by IBM/NSA), LOKI (Brown, Pieprzyk, and Seberry), and FEAL (Shimizu and Miyaguchi). A class of product ciphers called *Feistel ciphers* operates on half of the ciphertext at each round, then swaps the ciphertext halves after each round. Examples of Feistel ciphers include Lucifer and DES, both of which are commercial systems, the subject of the next section of this chapter.

7.3.4 Data Encryption Standard. Traditionally, the primary markets for code makers and computer makers have been the same: governments and banks. After World War II, computers were developed for both military and commercial purposes. By the mid-1960s, the leading computer maker was IBM, which could see that the growing role of electronic communications in commerce would create a huge market for reliable encryption methods. Over a period of years, mathematicians and computer scientists, including Horst Feistel at the IBM research lab in Yorktown Heights, New York, developed a cipher called Lucifer that was sold to Lloyds of London in 1971 for use in a cash-dispensing system.[36]

The U.S. National Security Agency (NSA) was in close touch with the Lucifer project, making regular visits to the lab (the constant flow of personnel between the NSA, IBM, and the mathematics departments of the major American universities tended to ensure that all new developments in the field were closely monitored). At roughly the same time, the National Bureau of Standards (NBS) was developing standard security specifications for computers used by the federal government. In 1973, the NBS invited companies to submit candidates for an encryption algorithm to be adopted by the government for the storage and transmission of unclassified information. (The government handles a lot of information that is sensitive but not sufficiently relevant to national security to warrant classification.)

IBM submitted a variation of its Lucifer cipher to the NBS, and after extensive testing by the NSA, this cipher was adopted as the nation's Data Encryption Standard (DES). The acronym actually refers to a document published as Federal Information Processing Standards Publication 46, or FIPS PUB 46 for short. This was published on January 15, 1977, and DES became mandatory for all "federal departments and agencies, for any . . . nonnational-security data."[37] The federal mandate also stated that commercial and private organizations were to be encouraged to use DES.[38] As a result, DES became widely used, especially in the banking industry.[39] The heart of DES is the Data Encryption Algorithm (DEA), which is described in a publication of the American National Standards Institute, titled *American National Standard for Information Systems—Data Encryption Algorithm—Modes of Operation, 1983,* referred to as ANSI X3.106-1983.

7.3.5 DES Strength. DES became, and remained, the de facto standard for commercial encryption until the late 1990s, when doubts about its strength relative to the rapid advances in computer hardware and software led to a quest for an eventual replacement. However, DES is still widely deployed, so more detailed discussion of its use is needed before discussing its replacement. The first thing to note is that the only

known method of deciphering data encrypted with DES without knowledge of the key is the use of brute force. This involves the computerized comparison of plaintext data with encrypted versions of the same data, using every possible key until both versions of the data match. With DES, the number of possible combinations is about 70 quadrillion. That is a very big number, and trying all those combinations within anything less than years requires relatively expensive hardware (or the carefully orchestrated application of large amounts of cheap hardware).

Technically speaking, the DEA is a combined substitution/transposition cipher, a product cipher that operates on blocks of data 64 bits, or 8 bytes, in length. Using 56 bits for the key produces a keyspace of 2^{56}, or 72,057,594,037,927,940, a number in the region of 70 quadrillion. A diagram of DES is shown in Exhibit 7.10.

The difficulty of attacking DES can be increased fairly easily if double or triple encryption is used, but despite this, there has always been something of a cloud over DES. At the time the DEA was approved, two Stanford University professors who are preeminent in twentieth-century cryptography, Martin Diffie and Whitfield Hellman, pointed out that the algorithm, as approved by the NBS, would be increasingly vulnerable to attack as computer equipment increased in power and came down in cost.

7.3.6 DES Weakness. As the author George Sassoon writes, "Although both the U.S. Department of Commerce and IBM deny it vigorously, everyone in the know insists that the NSA enforced a halving of the DES key length to ensure that they themselves could break the ciphers even if nobody else could." Although the NBS dismissed such criticisms, and the NSA flatly denied that they were behind any attempts to weaken the cipher, this opinion received some support from the NSA in 1986 when the agency announced it would no longer certify the DEA for nonclassified use, less than 10 years after the DES was approved. This move was prompted by the rapid development of parallel computers, which achieve amazing processing capabilities by using hundreds or even thousands of multiple processors, working in parallel. These machines offer enormous power at considerably less cost than traditional supercomputers. Perhaps the NSA could see the inevitability of something like the EFF DES Cracker, which was built in 1998 for less than $250,000 and broke a DES-encrypted message in fewer than three days.

The original Lucifer cipher used data blocks of 128 bits and a key of 112 bits. If this had been adhered to in the DEA, the difference in the number of possible key combinations would have been staggering. Although 2^{56}, the current keyspace, is a number greater than 7 with 16 zeroes behind it, 2^{112} is greater than 5 with 33 zeroes behind it. The practical consequence of this weakness in the DEA meant that the demand for stronger algorithms remained, and promising new ones emerged, such as Bruce Schneier's Blowfish.

There are still some positive aspects to DES that make it viable for some commercial uses. As was mentioned earlier, the cryptographic weakness of DES can easily be strengthened by double encryption, which doubles the difficulty of decryption, taking the task well into the realm of supercomputers and purpose-built, massively parallel machines. The fact that DES has been a standard for so long means that DES now is available in many forms, such as single-chip implementations that can be inserted into ROM sockets and integrated into all manner of hardware, such as expansion cards, PCMCIA cards, and smart cards.

7.4 PUBLIC KEY ENCRYPTION. Even with a longer key, the DEA still would have a major weakness, one that it shares with all of the other private key encryption

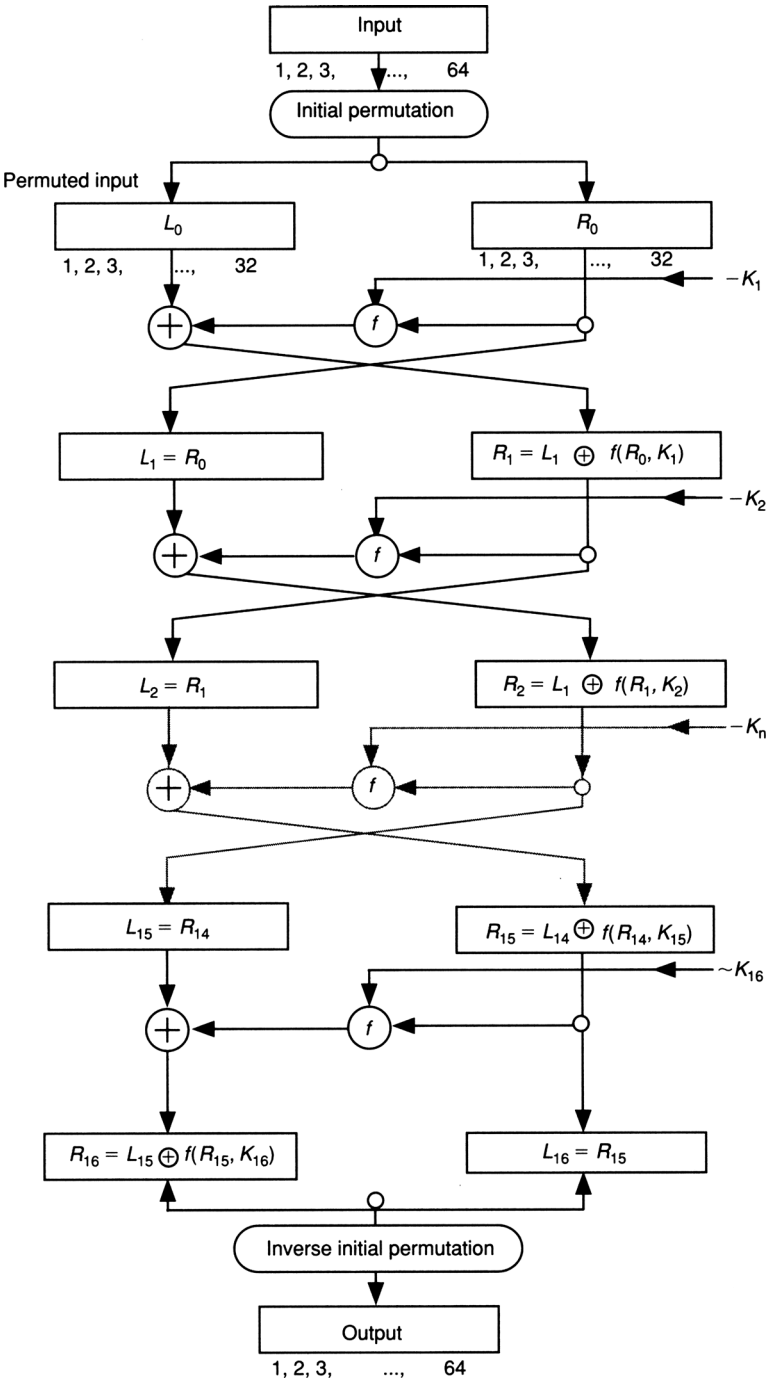


EXHIBIT 7.10 Diagram of DES

systems mentioned so far. That weakness is the need to keep the key secret. In this section we examine this problem, and the "public key" solutions that are now available.

7.4.1 Key-Exchange Problem. When password-protected data are sent from one place to another, either electronically or by hand, the need to transmit the password to the recipient presents serious obstacles. In cryptography, these are known collectively as the *key-exchange problem.* This is the way it is described by the Crypt Cabal[40]:

> If you want your friends to be able to send secret messages to you, you have to make sure nobody other than them sees the key. . . . [This is] one of the most vexing problems of all prior cryptography: the necessity of establishing a secure channel for the exchange of the key. To establish a secure channel, one uses cryptography, but private-key cryptography requires a secure channel!

So, even when using very powerful private key systems, such as DES, password or key distribution is a major problem. After all, the reason for encrypting valuable information in the first place is because it is assumed someone is trying to steal it or tamper with it. This implies a motivated and skilled adversary. Such an adversary is likely to use every opportunity to discover the password that will unlock the information. The password is perhaps most at risk from such an adversary when it is passed from one person to another. Although it sounds like the stuff of Bond movies, it actually is a very real and practical problem that had to be faced in many areas of legitimate organized activity, from businesses to public institutions, even when a powerful DEA-based computerized encryption system became available.

Suppose an encrypted file of sensitive accounting data needs to get to the head office. How does the recipient know the password needed to access the file? The sender could make a phone call. But will it be overheard? How is the identity of the person at the other end to be verified? A courier could be dispatched with a sealed envelope. The password could be encrypted. But all of these channels present problems. How to guarantee that the courier is honest or that the envelope will arrive intact? And if the password is encrypted, it will need a password itself, which will have to be transmitted. The recipient of the file can be provided with the password before the message is encrypted, but this is no guarantee that the password will not be intercepted. There are ways of making matters more difficult for the attacker, but the ideal solution would be to use a key that was useless to the attacker. This possibility is diagrammed in Exhibit 7.11.

7.4.2 Public Key Systems. A public key encryption system offers encryption that does not depend on the decryption key remaining a secret. It also allows the receiver of keys and messages to verify the source. The first published description of a public key cryptosystem appeared in 1976, authored by Stanford University professor Martin Hellman and researcher Whitfield Diffie. Ralph Merkle independently arrived at a similar system.

Ralph Merkle first proposed the idea of public key cryptography in 1974, and Martin Hellman and Whitfield Diffie brought the same idea to the public forum in 1976.[41] The idea was considered a seminal breakthrough, "for it had not occurred to anyone else in the long history of cryptology that the deciphering key could be anything other than the inverse of the enciphering key."[42] The Diffie-Hellman system employs a form of mathematics known as modular arithmetic. "Modular arithmetic is a way of restricting the outcome of basic mathematical operations to a set of integers with

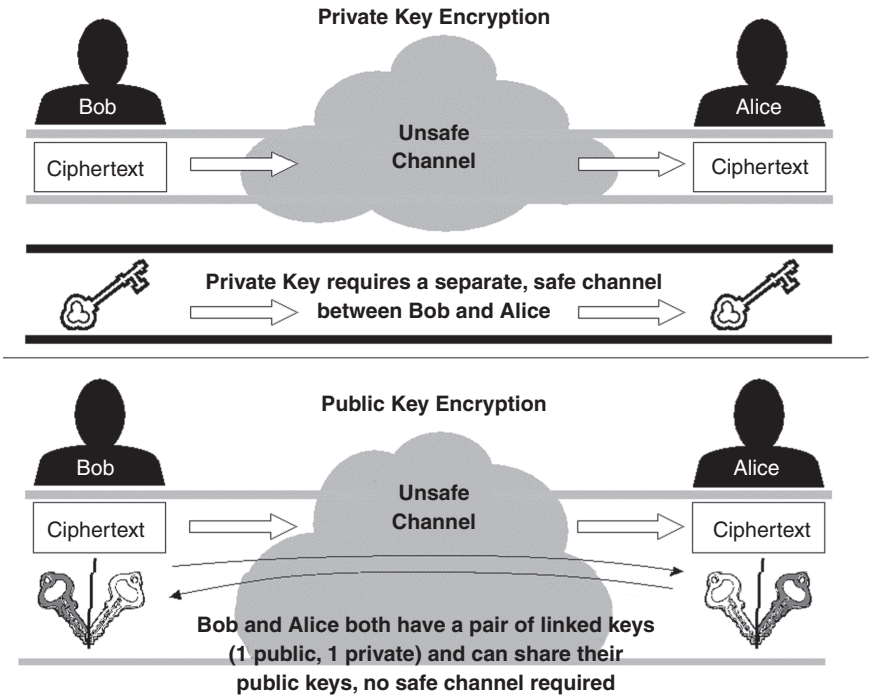


EXHIBIT 7.11 Comparison of Private and Public Key Encryption

an upper bound."[43] An excellent example of this mathematical principle is found by examining a military clock:

> Consider a clock on military time, by which hours are measured only in the range from zero to 23, with zero corresponding to midnight and 23 to 11 o'clock at night. In this system, an advance of 25 hours on 3 o'clock brings us not to 28 o'clock, but full circle to 4 o'clock (because 25 + 3 = 28 and 28 – 24 = 4). In this case, the number 24, an upper bound on operations involving the measurement of hours, is referred to as a modulus. When a calculation involving hours on a clock yields a large number, we subtract the number 24 until we obtain an integer between 0 and 23, a process known as modular reduction. This idea can be extended to moduli of different sizes.[44]

The Diffie-Hellman protocol allows two users to exchange a symmetric key over an unsecure medium without having any prior shared secrets. The protocol has two publicly known and widely distributed system parameters: p, a large prime integer that is 1,024 bits in length,[45] and g, an integer less than p. The two users wishing to communicate are referred to as Alice and Bob for simplicity's sake. They proceed in this way.

First, Alice generates a random private value a, and Bob generates a random private value b. Both a and b are [less than p]. Then they derive their public values using parameters p and g and their private values. Alice's public value is $g^a \bmod p$ and Bob's public value is $g^b \bmod p$. They then exchange their public values. Finally, Alice computes $g^{ab} = (g^b)^a \bmod p$, and Bob computes $g^{ba} = (g^a)^b \bmod p$. Since $g^{ab} = g^{ba} = k$, Alice and Bob now have a shared secret key k.[46]

This protocol introduced a concept to cryptography known as the discrete log problem. "The discrete log problem is stated as follows: given g, p, and $g^x \bmod p$,

what is x?"[47] It is generally accepted throughout the mathematical and cryptologic communities that the discrete log problem is difficult to solve, difficult enough for algorithms to rely on it for security.[48]

An algorithm to perform public key encryption was published in 1977 by Ronald Rivest of MIT, Adi Shamir of the Weizmann Institute in Israel, and Leonard Adleman of the University of Southern California. These three men formed the RSA Data Security Company, which was granted an exclusive license to the patent that MIT obtained on their algorithm. A large number of companies licensed software based on this algorithm, from AT&T to IBM and Microsoft. The RSA algorithm is currently at work in everything from online shopping to cell phones. Because it resolved the secret key dilemma, public key cryptography was hailed by many as a revolutionary technology, "representing a breakthrough that makes routine communication encryption practical and potentially ubiquitous," according to the Sci.Crypt FAQ, which states:

> In a public-key cryptosystem, E_K can be easily computed from some public key X, which in turn is computed from K. X is published, so that anyone can encrypt messages. If decryption D_K cannot be easily computed from public key X without knowledge of private key K, but readily with knowledge of K, then only the person who generated K can decrypt messages.[49]

The mathematical principles that make this possible are beyond the scope of this chapter. Somewhat more detail can be found in the RSA Laboratories' "Frequently Asked Questions About Today's Cryptography," which is distributed by RSA Data Security, the company that markets products based on the RSA algorithm. In brief, public key encryption is possible because some calculations are difficult to reverse, something pointed out by Diffie and Hellman, who first published the idea of public key encryption. Here is how RSA describes the calculations that make it possible (with minor clarification from the author):

> Suppose Alice wants to send a private message, m, to Bob. Alice creates the ciphertext c by exponentiating:
>
> $$c = m^e \mod n$$
>
> where e and n are Bob's public key. To decrypt, Bob also exponentiates:
>
> $$m = c^d \mod n$$
>
> where d is Bob's private key. Bob recovers the original message, m; the relationship between e and d ensures that Bob correctly recovers m. Because only Bob knows d, only Bob can decrypt.

This is diagrammed in Exhibit 7.12, which follows the scenario described. The lower part of the diagram uses numbers taken from an example given by Stallings. These numbers are much smaller than the actual numbers used by RSA. The point is that, given the ciphertext (c) and the public key (e,n) and knowledge of the algorithm, it is still impractical to decipher the message (m). This is because n is created by multiplying two prime numbers (normally represented as p and q) and e is derived from n combined with the secret key, d. To break the cipher, you need to factor a large number into a pair of prime numbers. How large? More than 150 digits in length (that is digits, not bits).

This cryptanalysis is very hard to do in a meaningful period of time, even with a very powerful computer. Large networks of computers have successfully factored

1. Private key, chosen Select two prime numbers, p and q	$p = 7$ and $q = 17$
2. Public key, calculated Calculate $n = pq$	$7 \times 17 = 119$
3. Public key, chosen Calculate $\phi(n) = (p - 1)(q - 1) = 96$ Select e, such that e is relatively prime to $\phi(n)$ and $< \phi(n)$	$e = 5$
4. Private key, calculated Determine d, such that de = 1 mod 96 and d < 96 Because $77 \times 5 = 385 = 4 \times 96 + 1$	$d = 77$
Result; Public key, KU = 5,119 Private key, KR = 77,119	

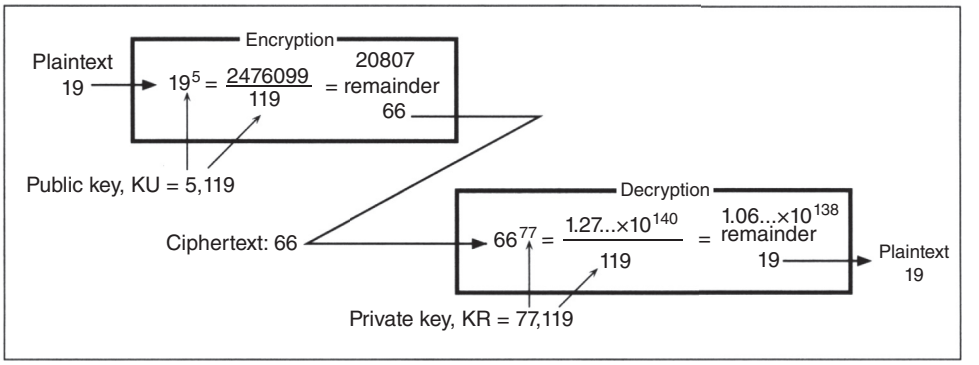


Exhibit 7.12 Public Key Diagram

a 100-digit number into two primes, but the RSA algorithm can use numbers even bigger if computer power and factoring algorithms start to catch up to the current implementations.

7.4.3 Authenticity and Trust. The point of the public key cryptosystems is to provide a means of encrypting information that is not compromised by the distribution of passwords, but public key encryption does not solve all problems associated with key exchange. Because the keys are considered public knowledge, some means "must be developed to testify to authenticity, because possession of keys alone (sufficient to encrypt intelligible messages) is no evidence of a particular unique identity of the sender," according to Sci.Crypt FAQ.[50]

This has led to key-distribution mechanisms that assure listed keys are actually those of the given entities. Such mechanisms rely on a *trusted authority,* which may not actually generate keys but does employ some mechanism which guarantees that "the lists of keys and associated identities kept and advertised for reference by senders and receivers are 'correct.' "[51] Another approach has been popularized by the program called Pretty Good Privacy, or PGP. This is the "Web of trust" approach that relies on users to distribute and track each other's keys and trust in an informal, distributed fashion.

Here is how RSA can be used to send evidence of the sender's identity in addition to an encrypted message. First, some information is encrypted with the sender's private key. This is called the *signature* and is included in the message sent under the public key encryption to the receiver. The receiver can "use the RSA algorithm *in reverse* to

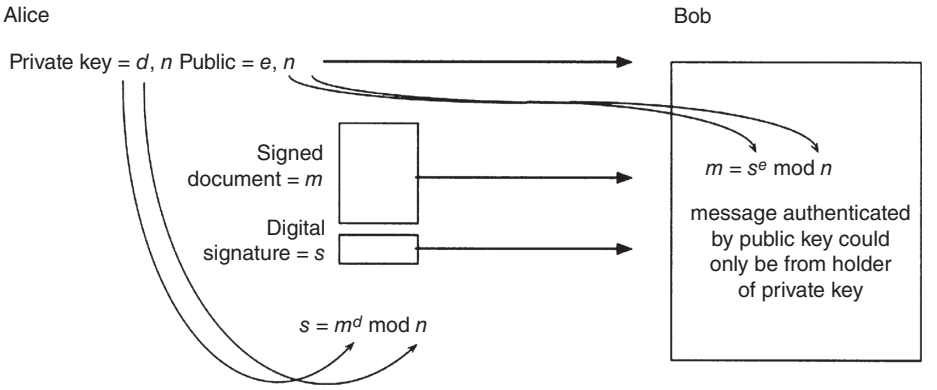


EXHIBIT 7.13 Authentication with RSA

verify that the information decrypts sensibly, such that only the given entity could have encrypted the plaintext by use of the secret key."[52]

What does "decrypts sensibly" mean? The answer involves something called a *message digest,* which is "a unique mathematical 'summary' of the secret message."[53] In theory, only the sender of the message could generate his or her valid signature for that message, thereby authenticating it for the receiver. Here is how RSA describes authentication, as diagrammed in Exhibit 7.13.

Suppose Alice wants to send a signed document m to Bob. Alice creates a digital signature s by exponentiating: $s = m^d \bmod n$, where d and n belong to Alice's key pair. She sends s and m to Bob. To verify the signature, Bob exponentiates and checks that the message m is recovered: $m = s^e \bmod n$, where e and n belong to Alice's public key.

7.4.4 Limitations and Combinations. As mentioned earlier, many products use RSA today, including Microsoft Windows, Lotus Notes, Adobe Acrobat, Netscape Navigator, Internet Explorer, and many more. In most of these examples, RSA is used for its authentication capabilities rather than for large-scale data encryption. That is because public key systems have one very noticeable downside: They are slow. This is balanced by the fact that they are harder to break. According to RSA, DES generally is at least 100 times as fast as RSA when implemented in software. In hardware, DES is between 1,000 and 10,000 times as fast, depending on the implementations. RSA may narrow the gap in coming years as more specialized chips are developed. However, public key algorithms are unlikely to ever match the performance of private key ciphers such as DES. Fortunately, there is a simple solution: Use a fast private key algorithm for the data encryption, but use a public key system to handle the key exchange and authentication, as diagrammed in Exhibit 7.14.

The private key encryption system might be DES or a system such as RC2 and RC4, both of which are available from RSA Data Security, or Schneier's Blowfish, which is freely available. Just as there are other private key systems besides DES, there are other public systems besides RSA. One method, called SEEK, is patented, trademarked, and marketed by Cylink of Sunnyvale, California. This method uses an alternative algorithm for public key distribution. Cylink manufactures a range of DES encryptors that use SEEK for key distribution.

7.5 PRACTICAL ENCRYPTION. The primary market for encryption systems and devices is communications. However, the development of Internet commerce has

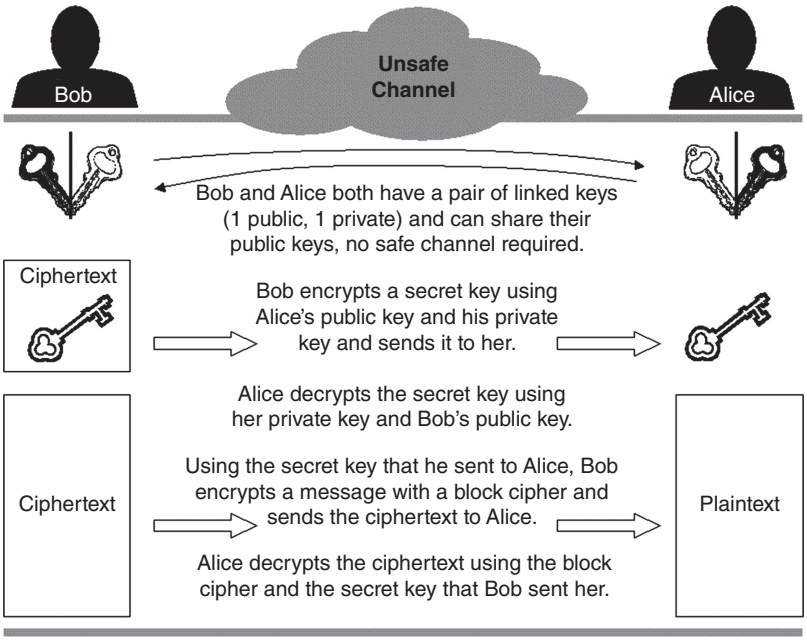


EXHIBIT 7.14 Combining Public and Private Key Encryption

resulted in a number of new and interesting crypto components that have considerable value for computer security.

7.5.1 Communications and Storage. If you look at the commercial products on the National Institute of Standards and Technology (NIST)'s list of approved DES implementations, most are designed to protect information when it is being communicated, not when it is sitting on a machine for local use. This is understandable when you look at the development of computing, which has spread outward from "fortress mainframe." Centralized data storage facilities lend themselves to physical access control. Encrypting data that stays behind walls and locked doors may be overkill in that scenario, particularly when there is a performance penalty involved.

Encryption was reserved for data in transit, between computers, across wires. This philosophy was extended to file servers on networks. File encryption on the server was not considered a priority as people assumed the server would be protected. Data encryption on stand-alone machines and removable media is a relatively recent development, particularly as more and more confidential data are packed into physically smaller and smaller devices. There are now many products with which to implement file encryption.

7.5.2 Securing the Transport Layer. One of the most visible examples of encryption at work in computer security today is the security icon people see in their Web browser; see Exhibit 7.15 for examples of Netscape Navigator and Microsoft Internet Explorer. This is an example of something called transport layer security, which uses protocols that go by the name of SSL and TLS.

7.5.2.1 Popular Protocols. SSL stand for Secure Sockets Layer, the software encryption protocol developed by Netscape and originally implemented in Netscape

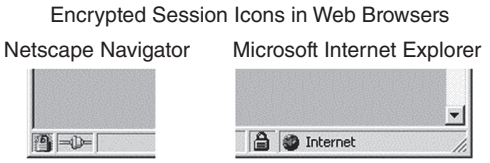


EXHIBIT 7.15 SSL 3.0 in Action

Secure Server and the Netscape Navigator browser. SSL is also supported by Microsoft Internet Explorer and a number of other products. TLS stands for Transport Layer Security, the name given to an Internet standard based on SSL, by the IETF (as in Internet Engineering Task Force, RFC 2246). There are minor differences between SSLv3.0 and TLSv1.0 but no significant differences as far as security strength is concerned, and both protocols interoperate with each other.

The TLS is a protocol, a standardized procedure for regulating data transmission between computers. It is actually composed of two layers of protocol. At the lowest level is the TLS Record Protocol, which is layered on top of some reliable transport protocol, typically the TCP in TCP/IP, the set of protocols that run the Internet. The TLS Record Protocol provides connection security that is both private (using symmetric cryptography for data encryption) and reliable (using a message integrity check). Above the TLS Record Protocol, encapsulated by it, is the TLS Handshake Protocol. This allows the server and client to authenticate each other, a major role for TLS in various forms of e-commerce, such as Internet banking. The TLS Handshake Protocol can also negotiate an encryption algorithm and cryptographic keys before any application protocol sitting on top of it, such as HTTP, transmits or receives its first byte of data (see Exhibit 7.16).

7.5.2.2 Properties of TLS. In providing connection security, the TLS Handshake Protocol delivers three basic properties. The identity of the parties can be

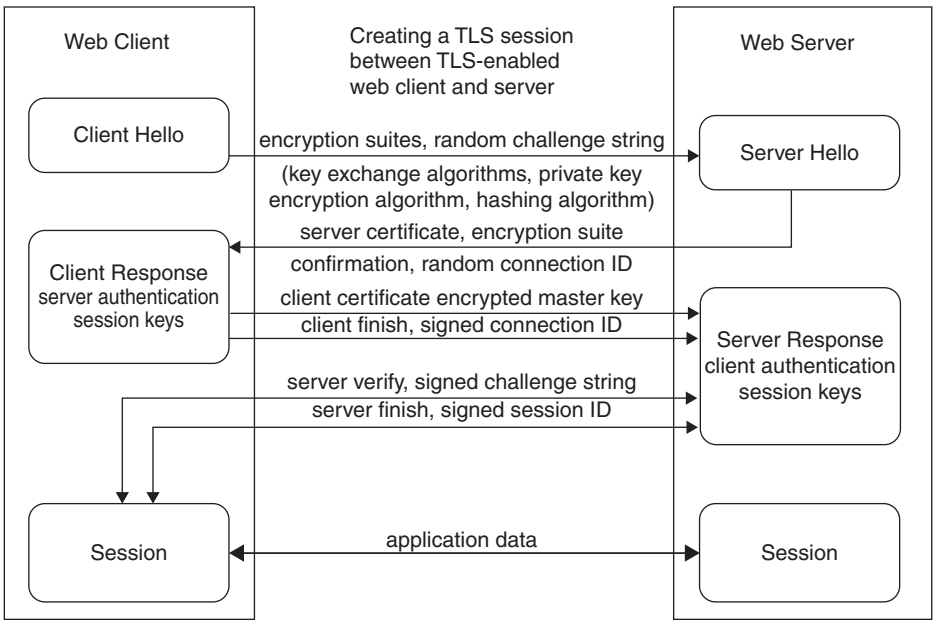


EXHIBIT 7.16 Creating a TLS Session

authenticated using public key cryptography (such as RSA). This authentication can be made optional, but typically it is required for at least one of the parties (e.g., the Yahoo! Travel server authenticates itself to the user's browser client, but the user's client does not authenticate itself to the Yahoo! Travel server, a distinction discussed in a moment).

The second and third basic properties of the TLS Handshake Protocol are that a shared secret can be securely negotiated, unavailable to eavesdroppers, even by an attacker who can place itself in the middle of the connection; and the protocol's negotiation is reliable. In the words of RFC 2246: "no attacker can modify the negotiation communication without being detected by the parties to the communication."

TLS can use a variety of encryption algorithms. For the symmetric encryption that is part of the Record protocol, DES or RC4 can be used. The keys for this symmetric encryption are generated uniquely for each connection and are based on a secret negotiated by another protocol (such as the TLS Handshake Protocol). The record protocol includes a message integrity check using a keyed MAC, with secure hash functions such as SHA and MD5, used for MAC computations. The encryption suite to be used for a specific connection is specified during the initial exchange between client and server, as shown in Exhibit 7.16.

7.5.2.3 Tested in the Real World. TLS/SSL has been widely used and extensively tested in the real world, and thoroughly probed by real cryptographers. Some of the caveats and limitations noted by these and other experts follow. The first is that neither a good standard nor a good design can guarantee a good implementation. For example, if TLS is implemented with a weak random number seed, or a random number generator that is not sufficiently random, the theoretical strength of the design will do nothing to protect the data that are thus exposed to potential compromise. (Although beyond the scope of this chapter, Pseudo-Random Number Generators, or PRNGs, play a vital part in many cryptographic operations, and they are surprisingly difficult to create; unless they closely simulate true randomness, an attacker will be able to predict the numbers they generate, thus defeating any scheme that relies on their "random" quality.)

The second major caveat is that, if clients do not have digital certificates, the client side of the TLS session is not authenticated. This presents numerous problems. Most of today's "secure" Web transactions, from airline tickets booked at Yahoo Travel to shares traded at most online brokerages, represent a calculated risk on the part of the vendor. Although the client doing the buying is assured, by means of the merchant certificate, that the merchant at www.amazon.com really is Amazon, the merchant has no digital assurance that the client computer belongs to, or is being operated by, the person making the purchase. Of course, there are other assurances, such as the match between the credit card that the purchaser supplies and the other personal details that go along with it, such as billing address. But the merchant is still risking a charge-back and possibly other penalties for a fraudulent transaction.

In the case of larger and more sensitive financial transactions, the need to be assured of the client's identity is greater. A digital certificate is a step in the right direction, but it is a step many merchants have not yet taken, for several reasons. The first is the cost of issuing certificates to customers, and the second is the difficulty of getting those certificates onto their systems. Some merchants have decided that the cost and effort are worth it. For example, the Royal Bank of Scotland took this approach with its online banking system back in 1998.

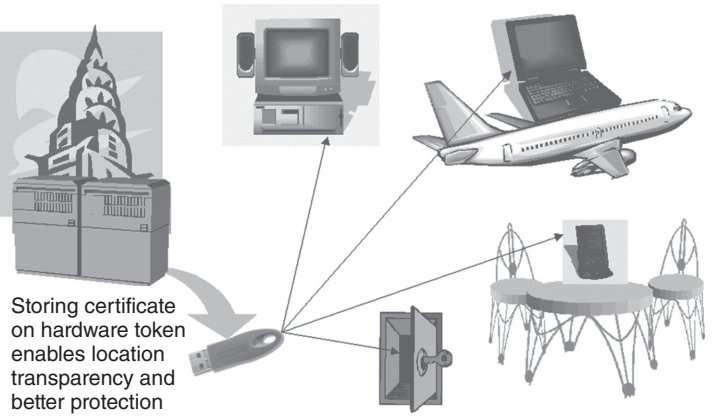


EXHIBIT 7.17 Using a Hardware Token for Digital Signatures

There are other issues. The user needs to protect the certificate, even from such threats as hardware failure (user reformats the drive, loses the certificate) or unauthorized use (a family member uses the computer and thus has access to the certificate). Furthermore, the user needs to be able to move the certificate, for example, onto a laptop computer so that the bank account can be accessed while traveling. The obvious answer is to place the certificate on a robust removable medium (see Exhibit 7.17). Such media are generically referred to as hardware tokens. A standard for tokens has not yet emerged. Smart cards are an obvious choice, but card readers need to be deployed. There are alternatives, such as putting the certificate on a floppy disk or on a small key fob that plugs into a USB port.

7.5.2.4 Cost of Secured Transactions. For companies looking to perform highly secure transactions today, using SSL without client-side authentication is proving acceptable in the short term, at least for some categories of transaction. Even then it can be costly, in terms of either dollars or processing power. Although TLS is an open standard, and Netscape has provided crucial parts of the technology royalty free, there is still the question of which algorithms to use. Some algorithms are more expensive than others, and not always in obvious ways. For example, you have to license RC4, whereas DES is free, but RC4 is optimized for a 32-bit processor and DES is not.

Furthermore, research shows that the amount of "hits" that a Web server can handle drops dramatically when those hits require TLS (and it drops a whole lot more when processing client authentication as well as server authentication). The answer here may be specialized hardware. Several companies, such as IBM and Rainbow Technologies, make crypto-coprocessor cards that relieve the server's CPU of the specialized math processing involved in crypto. They are cheaper than adding another server to keep up with the very demanding task of providing secure Web transactions.

7.5.3 X.509v3 Certificate Format. Another example of encryption widely used in computer security today is X.509. This is not a rocket ship but a standard for digital certificates, described earlier in this chapter. The International Telecommunication Union's Telecommunication Standardization Sector (ITU-T)'s X.509 standards document states: "Virtually all security services are dependent upon the identities of the communicating parties being reliably known, i.e. authentication." Consider how this affects Web transactions. The preceding section described how SSL can encrypt



EXHIBIT 7.18 Digital Certificate

Web pages sent from Web server to Web client, and vice versa, but it cannot assure the identity of the parties involved. The X.509 standard helps to address this problem, which negatively impacts the profitability of Web-based businesses.

When a Web user asks for assurance that the bn.com Web site is actually Barnes & Noble, it can be provided by way of a digital certificate (see Exhibit 7.18). This means that an entity, known as a certificate authority (CA), has taken considerable pains to reliably identify, and consequently certify, the merchant as the rightful owner of an encryption key. This key is the public half of a uniquely and mathematically related public/private key pair, such that a message encrypted with the public key can only be decrypted with the corresponding private key.

Individuals, as well as merchants, can have a public/private key pair. A bank might then access that public key, and use it, plus the bank's private key, to encrypt the account details it sends to customers over the Web. Only the customer with the right private key can decrypt this information, using the bank's public key. At the same time, customers know the statement information can only have come from the bank (otherwise the bank's public key would not work to decrypt it). Customers also know,

EXHIBIT 7.19 X.509 Certificate Format

Version	Identifies the Certificate Format
Certificate Serial Number	Number that is unique within the issuing CA
Signature Algorithm Identifier	Identifies the algorithm used to sign the certificate, together with any necessary parameters
Issuer	X.500 name of the issuing CA
Validity Period	Pair of dates between which the certificate is valid
Subject	X.500 name of the holder of the private key corresponding to the public key certified by the certificate
Subject Public Key Information	Public key for the subject, plus an identifier for the algorithm with which this public key is to be used

thanks to an encrypted message digest (a digital fingerprint of the message contents), that the data they get from the bank has not been altered. Thus, it is very difficult for either party to claim that it never took place. In this way, digital certificates can enhance confidentiality, integrity, and nonrepudiation.

7.5.3.1 ISO/IEC/ITU 9594-8 a.k.a. X.509. The management of public keys is the task of Public Key Infrastructure (PKI), of which the X.509 standard is an important part. For example, an organization's employees can perform secure business communications over the Internet, such as contract negotiation, using PKI. To engage in a secure transaction with someone, it is necessary to find and access the other person's public key, and vice versa. The answer is to publish public keys in the form of a digital certificate, then use some form of directory to locate them. In order for different systems to interoperate, standards for directories have been developed, notably X.500. This standard applies such elements of directory standardization as a hierarchical naming convention:

> Country, Organization, Common Name.
> So Fred Jones of Megabank might have the X.500 name:
> [Country = US, Organization = Megabank, Inc., Common Name = Fred Jones]

A means of locating digital certificates to verify identities was a logical extension of the standard, thus X.509 was developed, officially known as ITU-T X.509 (formerly CCITT X.509) and also ISO/IEC/ITU 9594-8. In X.509 there is a definition of a basic certificate format, which consists of seven fields shown in Exhibit 7.19.

The certificate format has evolved considerably since 1988. The original format is now referred to as X.509v1. When X.500 itself was revised in 1993, two more fields were added to support directory access control, resulting in the X.509v2 format.[54] X.509v2 added unique identifiers for the issuer and the subject, optional bit strings used to make the issuer and subject names unambiguous in the event that the same name is later reassigned to different entities. Suppose that Fred Jones, whose assigned X.500 name was given earlier, is an executive vice president of Megabank, but is then hired away by a competitor. Megabank deassigns his name, but if a different Fred Jones, a programmer, then comes to work for Megabank, he is effectively reassigned the same X.500 name:

> [Country = US, Organization = Megabank, Inc., Common Name = Fred Jones]

This poses authorization problems for any access control lists attached to X.500 data objects, due to the difficulty of identifying all of the access control lists that grant privileges to a particular user's name. The unique identifier field added in X.509v2 provides somewhere to put a new value whenever a name is reused. In fact, a better solution is to use a better distinguisher in the X.500 name, such as:

[Common Name = Fred Jones, Employee Number = 1000002]

In 1993, when the Internet Privacy Enhanced Mail (PEM) RFCs were published, they included specifications for a public key infrastructure based on X.509v1 certificates. Attempts to deploy PEM, however, revealed deficiencies in the Version 1 and 2 certificate formats. Consequently, ISO/IEC/ITU and ANSI X9 developed the X.509v3 format, which greatly extends the capabilities of the format by providing extension fields and broader naming options in X.509v3.

7.5.3.2 Extending the Standard. Extensions were added in Version 3 to address problems discovered while implementing Version 1 and 2 certificates. These can be seen in the diagram in Exhibit 7.20. Particular extension field types may now be specified in standards or defined and registered by any organization or community. Each extension field is assigned a type by means of an object identifier, registered in the same way that an algorithm is registered. Although theoretically anyone can define an extension type, to achieve practical interoperability, common extension types need to be understood by different implementations. Thus, the most important extension types are standardized. But when X509v3 is used within a closed group—for example,

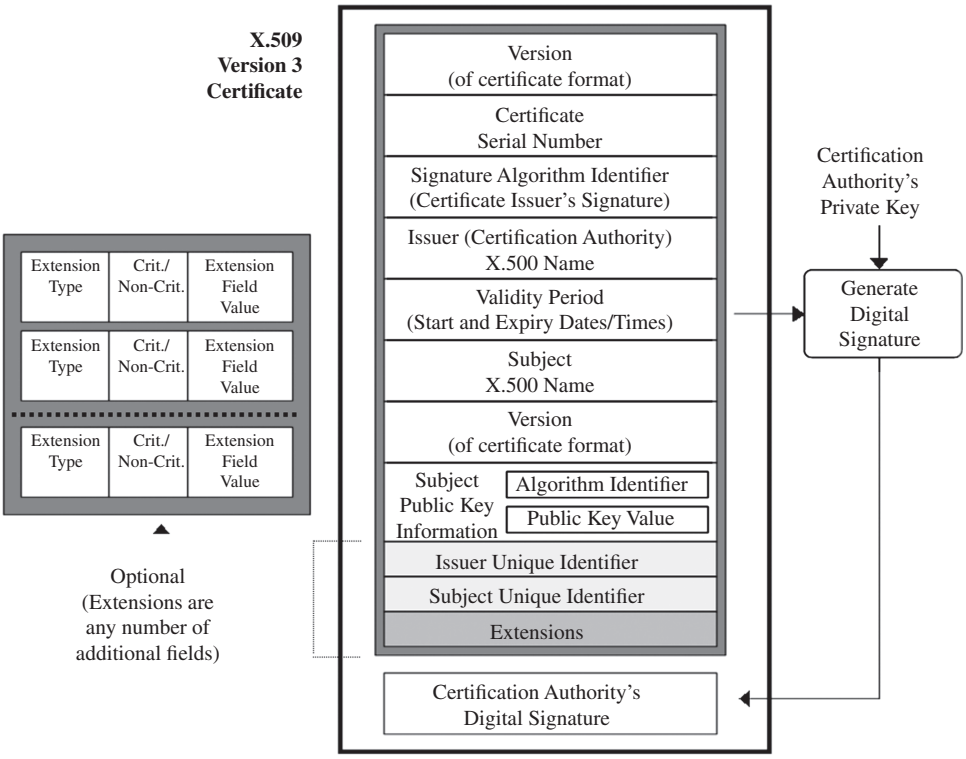


EXHIBIT 7.20 X.509v3 Certificate

a group of business partners—it is possible to define unique extension types to satisfy specific needs.

7.5.3.3 X.509 Sources, Issues, and CAs. Someone managing an e-commerce project does not necessarily need to know X.509 in detail but should at least read the Arsenault and Turner document (see Section 7.8, "Further Reading," at the end of the chapter); it clearly describes not only X.509 but the role it plays in PKI (which they define as "the set of hardware, software, people, policies and procedures needed to create, manage, store, distribute, and revoke certificates based on public-key cryptography"[55]). Also very helpful are the presentations by VeriSign's Warwick Ford, which NIST has online at its Web site. For the e-commerce developer who wants more detail, the next step is Ford's book, coauthored with fellow VeriSign executive Michael Baum, *Secure Electronic Commerce*.[56] This documents other important aspects of X.509, such as the Certificate Revocation List, used to revoke certificates before they expire (e.g., if the private key has been compromised). A copy of the standard, available online, at the International Telecommunication Union (ITU) Web site (www.itu.int), is also valuable.

The extensions and improvements in the X.509v3 certificate format greatly increase its usefulness, but providing a uniform method of going beyond the standard does raise the specter of a lack of standardization. This is something that the IETF's PKIX working group is addressing. And there are other issues to consider when evaluating X.509 as a security technology, many of which are raised by Ed Gerck of the Meta-Certificate Group. Articles at the group's Web site point out that X.509 does not address "the level of effort which is needed to validate the information in a certificate."[57] In other words, some security issues are beyond the scope of X.509, but they do need to be considered when deploying systems that rely on these certificates. For example, it does not make sense to rely on a digital certificate if the measures taken to assure the identity of the owner and user of the certificate are not commensurate with the risk involved in relying on the certificate. Furthermore, transactions that do not use certificates on both sides will remain inherently problematic.

These issues point to the importance of the role played by the CA. As mentioned earlier, CAs are the entities that issue and sign certificates. Each has a public key that is listed in the certificate. The CA is responsible for scheduling an expiration date and for revoking certificates when necessary. The CA maintains and publishes a Certificate Revocation List (CRL).

In other words, ensuring the validity of certificates entails a lot of maintenance. The CRL, for example, is crucial if certificates are compromised or found to be issued fraudulently. This happened in 2001 when a number of VeriSign certificates were found to be issued in error to someone posing as Microsoft. Because some computer users now rely on certificates to guarantee the authenticity of software upgrades and components, failure to check the revocation list before downloading certified code could result in malicious code attacks.

Problems with certificates have the potential for widespread impact because the authority in certificates is hierarchical, as shown in Exhibit 7.21. When a CA issues a certificate, it signs it with its own key. Anyone relying on certificates issued by that CA needs to know by what authority the CA is issuing that certificate. To simplify, there are two possible answers. The CA is self-certifying, that is, providing its own "root" key, or it is relying on another CA for the root key. Clearly, any compromise of the root key undermines all certificates that gain their authority from it.

See Chapter 37 of this *Handbook* for a more extensive discussion of PKI and CAs.

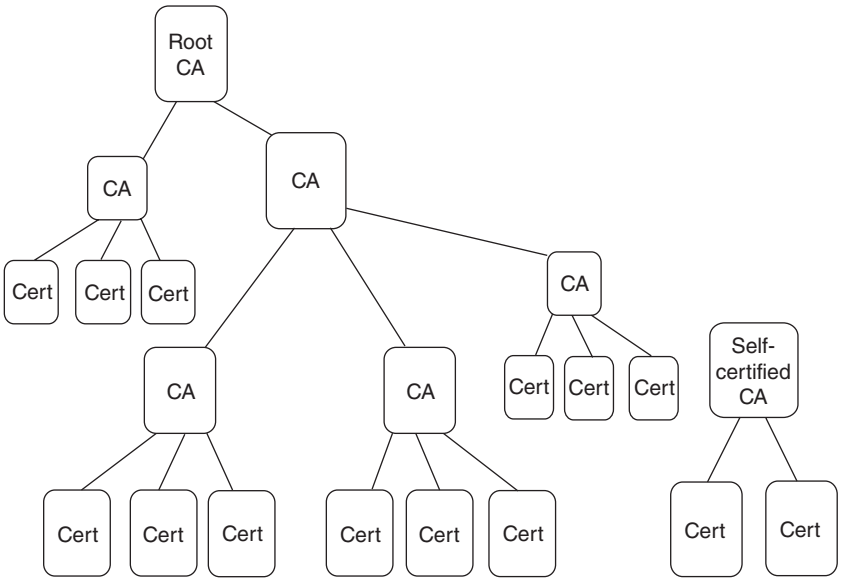


Exhibit 7.21 Certificate Authorities and the Root Key

7.6 BEYOND RSA AND DES. Cryptography research and development did not stop with the development of the RSA algorithms. Events in the last two decades of the twentieth century and the first decade of the twenty-first, and their implications, are discussed in this final section of the chapter, which concludes with some warnings on implementing encryption.

7.6.1 Elliptic Curve Cryptography. In 1985, Neal Koblitz from the University of Washington and Victor Miller of IBM independently discovered the application of elliptic curve systems to cryptography. When applied to public key cryptography, elliptic curve arithmetic has been found to offer certain advantages over first-generation public key techniques such as Diffie-Hellman and RSA.

The security of elliptic curve algorithms is based on the same principle as the Diffie-Hellman algorithm, the discrete log problem, as described in Section 7.4.2. The advantages to elliptic curve algorithms lie in the key size needed to achieve certain levels of security. As one scales security upward over time to meet the evolving threat posed by eavesdroppers and hackers with access to greater computing resources, elliptic curves begin to offer dramatic savings over the old, first-generation techniques.[58]

Until 2010, public key systems used 1,024 bits or 2,048 bits for creating keys. NIST recommended that after 2010, these systems be upgraded to a system that can provide adequate security. One way of doing this would be to increase the key size that is used. However, systems that are in place today become increasingly cumbersome the larger the key size. The NSA is endorsing elliptic curve cryptography, stating on its Web site that it has implemented elliptic curve public key cryptography systems to protect both classified and unclassified information.[59] Elliptic curve systems offer a way to increase key size moderately when more security is required. Exhibit 7.22 shows the NIST recommended key size that RSA or Diffie-Hellman should use to protect the transportation of symmetric keys of various sizes as well as the corresponding elliptic curve key size.

EXHIBIT 7.22 NIST Recommended Key Sizes

Symmetric Key Size (bits)	RSA and Diffie-Hellman Key Size (bits)	Elliptic Curve Key Size (bits)
80	1024	160
112	2048	224
128	3072	256
192	8192	384
256	15360	521

Source: National Security Agency, "The Case for Elliptic Curve Cryptography," www.nsa.gov/business/programs/elliptic_curve.shtml.

Thus, in order to use RSA to protect a 256-bit AES key, one should use a key of 15,360 bits, which is an order of magnitude larger than the key sizes currently in use throughout the Internet. However, an elliptic curve key would need to be only 521 bits. Elliptic curve algorithms can use smaller keys, because the math involved makes the inverse, or decryption, operations harder as the key length increases.[60]

Another feature that makes elliptic curves appealing is the fact that they are more efficient than the current implementations of public key cryptography, which tend to be relatively slow, causing them to be used more as key distribution methods than data encryption methods. Exhibit 7.23 shows the ratio of Diffie-Hellman computations versus elliptic curve computations for each of the key sizes listed in Exhibit 7.22.[61]

7.6.2 RSA Patent Expires. On September 6, 2000, RSA Security released the RSA public key encryption algorithm into the public domain. This means that anyone can now create products that incorporate this algorithm (provided it is their own implementation and not one licensed from RSA). In effect, RSA Security waived its rights to enforce the patent for any development activities that include the RSA algorithm occurring after September 6, 2000. The U.S. patent for the RSA algorithm actually expired on September 20, 2000. The result has been an even broader use of public key encryption, at lower cost.

The RSA patent was always somewhat controversial, because it applied to a piece of mathematics, which is not what most people think of when they think of an invention. The owners of the patent were never able to expand protection beyond the United States. As a result, versions of public key encryption based on alternatives to the RSA algorithm were developed and marketed outside the country, by companies

EXHIBIT 7.23 Relative Computation Costs of Diffie-Hellman and Elliptic Curves

Security Level (bits)	Ratio of DH Cost : EC Cost
80	3:1
112	6:1
128	10:1
192	32:1
256	64:1

Source: National Security Agency, "The Case for Elliptic Curve Cryptography," www.nsa.gov/business/programs/elliptic_curve.shtml.

like Ireland's Baltimore Technologies, Finland's F-Secure, and Israel's Algorithmic Research. Now encryption companies can dispense with the costly maintenance of multiple versions of their public key products (U.S. and non-U.S.). In addition, U.S. companies can develop and market RSA-based products. Large companies actually can "roll their own" public key encryption schemes for internal use, based on a proven, royalty-free algorithm.

7.6.3 DES Superseded. RSA Security, the company that tried to make the RSA algorithm synonymous with public key encryption, played a leading role in the other watershed crypto event of 2000, the naming of a successor to DES, the Data Encryption Standard. As noted earlier, projects like the EFF DES Cracker showed that a computer built for less than $250,000 could decipher a DES-encrypted message in fewer than three days. In fact, this was part of the "DES Challenges" sponsored by RSA Security. DES Challenge I was won by Rocke Verser of Loveland, Colorado, who led a group of Internet users in a distributed brute force attack. The project, code-named DESCHALL, began on March 13, 1997, and was successfully completed some 90 days later. DES Challenge II consisted of two contests posted on January 13 and July 13, 1998. The first contest was cracked by a distributed computing effort coordinated by distributed.net, which met the challenge in 39 days. The second contest was the one solved by EFF's purpose-built DES Cracker.

The effect of these projects was to focus attention on the need for stronger encryption. Companies and government agencies wanting to archive sensitive data need it to remain secure for decades, not days. However, as predicted in the 1970s, advances in computer power rendered "obsolete" the DEA, the widely used private key algorithm that forms the basis of the DES. Of course, the term "obsolete" is relative in this context. DES is not obsolete when applications need to encrypt bulk data to keep it confidential for a limited period of time, and a lot of data falls into this category. As Exhibit 7.24 shows, there is a direct relationship among time, technology, and the degree of protection that any ciphersystem provides.

In 1997, the U.S. Government began the process of establishing a more powerful standard than DES, known as Advanced Encryption Standard (AES). This is a Federal Information Processing Standard (FIPS) Publication, FIPS 197, specifying "a

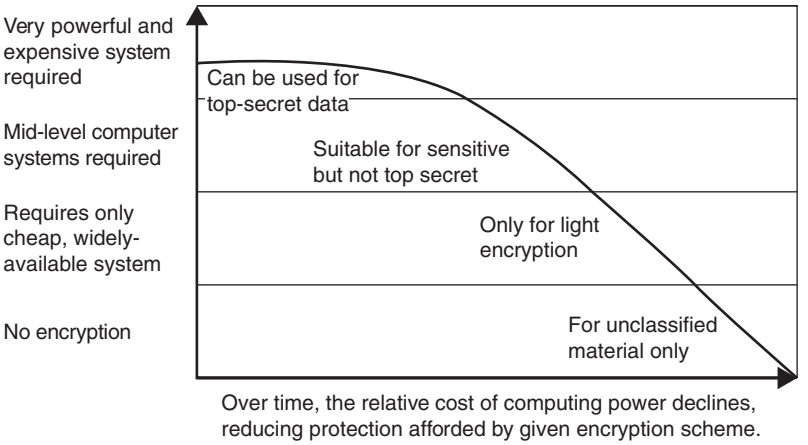


EXHIBIT 7.24 Relationship among Time, Technology, and Protection

cryptographic algorithm for use by U.S. Government organizations to protect sensitive (unclassified) information." The government anticipated correctly that AES would be "widely used on a voluntary basis by organizations, institutions, and individuals outside of the U.S. Government—and outside of the United States—in some cases."

In essence, a competition was held to find the best possible algorithm for the job, and the winner, chosen in October 2000, was Rijndael (pronounced "Rhine Doll"). This algorithm was developed specifically for the AES by two cryptographers from Belgium, Dr. Joan Daemen and Dr. Vincent Rijmen. Rijndael is a block cipher with a variable block length and key length. So far, keys with a length of 128, 192, or 256 bits have been specified to encrypt blocks with a length of 128, 192, or 256 bits. (All nine combinations of key length and block length are possible.) However, both block length and key length can be extended very easily in multiples of 32 bits. Rijndael can be implemented very efficiently in hardware, even on smart cards.

7.6.4 Quantum Cryptography.

A new basis for computation will profoundly affect cryptographic strength in the coming decades. This section provides a brief and nontechnical summary of the science of quantum computation and quantum cryptography.

7.6.4.1 *Historical Perspective.*

The entirety of this chapter has focused on the status of cryptography as it currently exists. The classic computer has been sufficient to perform the computations and processes required of AES, RSA, and all of the cryptographic systems and algorithms that have been explored since the advent of cryptography. Although modern computers are fundamentally the same as they were in the 1950s, the machines we use today are significantly faster.[62] Even though the speed has increased, the primary task of computers has remained the same: "to manipulate and interpret an encoding of binary bits into a useful computational result."[63] To push the bounds of computer performance ever forward, computer scientists' goal has "been the reduction of size in the transistors used in modern processors."[64]

Early computers were constructed of gates and storage "bits" made of many thousands of molecules. The components of today's processors are moving in the direction of a few hundred molecules. The computing industry has always known that miniaturization would reach a barrier below which circuits could not be built, because their fundamental physical behavior would change.[65]

The components of modern computers are reaching this barrier; should transistors become much smaller, they will "eventually reach a point where individual elements would be no larger than a few atoms."[66] Computer scientists are concerned about this continual shrinking, because at the atomic level, the laws of quantum mechanics will govern the properties and behavior of circuits, not the laws of classical mechanics.[67]

The science of quantum mechanics is not fully understood by scientists; it was initially thought to be a major limitation to the evolution of computer technology.[68] It was not until 1982 that the scientific community saw any benefit from the unusual effects associated with quantum mechanics. That year, Richard Feynman theorized about a new type of computer that would harness the effects of quantum mechanics and use these effects to its advantage.[69] In 1985, David Deutsch of the University of Oxford published a "ground breaking theoretical paper describing how any physical process could be modeled perfectly (in theory) using a quantum computing system."[70] He further argued that a quantum system would be able to execute tasks that no modern computer could perform, such as *true* random number generation.[71] "After

Deutsch published this paper, the search began to find interesting applications for such a machine."[72]

7.6.4.2 Fundamentals. A "quantum" is "the smallest amount of a physical quantity that can exist independently, especially a discrete quantity of electromagnetic radiation."[73] Quantum mechanics explains the physics and behaviors of particles, atoms, and energy.[74] The idea of a quantum computer is based on the phenomena that occur at the atomic and subatomic level, which are explained by quantum mechanics and defy all classical laws of physics.[75] These phenomena will be covered in more detail shortly; it is necessary at this point, however, to explain several fundamental differences between classical modern computers and the idea of a quantum computer.

Classical computers store and process information in units called bits, represented as a zero (0) or a one (1) in a computer's transistors. Bits are then organized into bytes, a series of eight bits. Thus, the information stored on a computer is stored as individual bits grouped into bytes. Therefore, a document "comprised of n-characters stored on the hard drive of a typical computer is accordingly described by a string of 8n zeros and ones."[76] It is important to emphasize that bits "can *only* exist in one of two distinct states, a '0' *or* a '1'."[77] This leads to the first difference between classical computers and quantum computers.

Quantum computers store and process information in units called quantum bits, referred to as "qubits." "Qubits represent atoms, ions, photons or electrons and their respective control devices that are working together to act as computer memory and a processor."[78] Similar to a classic bit, a qubit is represented as a 0 or a 1. Unlike a classic bit, a qubit can also exist in a *superposition* of both a 0 *and* a 1. In other words, it is possible for a single qubit to exist as a 0, a 1, or simultaneously as both a 0 *and* a 1. A qubit that is in two positions at once is said to be in its coherent state.[79] This can be explained more coherently with an example:

> If a coin is flipped in a darkened room, the result of the coin being flipped is mathematically just as likely to be heads or tails. While the light is off, the coin is in a superposition—whereby it is both heads and tails at once, because an [observer] cannot see which it is. If [the observer] turns on the light, [he or she] "collapses" the superposition, and forces the coin to be either heads or tails by measuring it. Measuring something destroys the superposition, forcing it into being in just one classical state.[80]

This coherent state leads to the phenomenon that would make a quantum computer exponentially more powerful than any computer to date; this is the phenomenon called "quantum parallelism."[81] Essentially, because a qubit in a coherent state holds two values at once, a single operation done on such a qubit would act on both values at the same time.[82] "Likewise, a two-qubit system would perform the operation on four values and a three-qubit system on eight [values]."[83] To summarize, an operation done on a system of n qubits would act on 2^n values simultaneously.[84] Exhibit 7.25 shows this concept using a system containing three qubits, which represent eight states simultaneously.

"The very property that makes quantum computing so powerful also makes it very fragile and hard to control."[85] In order to harness the power of quantum parallelism, scientists need to be able to read and measure the output from the operations performed on groups of qubits. Herein lies the problem of decoherence. When a qubit in the coherent state measurably interacts with the environment, it will immediately decohere and resume one of the two classical states, either a 0 or a 1, and it will no longer exhibit

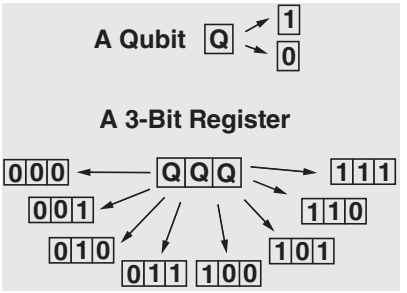


Exhibit 7.25 Three-Qubit System
Source: Simon Bone and Matias Castro, "A Brief History of Quantum Computing," Imperial College, London, www.doc.ic.ac.uk/~nd/surprise 97/journal/vol4/spb3/.

its dual-state ability. In other words, simply looking at a qubit can cause it to decohere, and this makes measuring qubits directly impossible.[86]

If scientists are unable to measure something directly, then they must find a way to measure indirectly, or a practical quantum computer will never be made. One possible answer lies in another property of quantum mechanics called entanglement. Entanglement is an obscure attribute that involves two or more atoms or particles. When certain conditions are met or certain forces are applied to two or more particles, then they can become entangled, whereby the particles exhibit opposite properties. The entangled particles will remain entangled, no matter the physical distance between them, and one entangled particle will always be able to communicate with its partner. Particles spin either up or down, and this spin is how scientists measure information about the particles. The property of coherence tells us that a particle will spin both up and down simultaneously until a scientist looks at it and measures it. "The spin state of the particle being measured is . . . communicated to the correlated particle, which simultaneously assumes the opposite spin direction to that of the measured particle."[87] Thus, entanglement could allow scientists to know the value of a qubit without actually looking at one. Scientists admit that entanglement is a difficult notion; they are still exploring the concept.[88] They also acknowledge that it could be years before a workable solution to the problem of measuring information in a quantum system is discovered.[89]

7.6.4.3 Impacts. Although quantum computers, in theory, can perform any task that a classical computer can, this does not necessarily mean that a quantum computer will always outperform a classical computer. Multiplication is an often-cited example of something that would be done just as quickly on a classical computer as on a quantum computer.[90] From the early stages of quantum computing, scientists knew that to demonstrate the superior computing power, new algorithms would have to be designed to exploit the phenomenon of quantum parallelism. Such algorithms are complex and difficult to devise, but two are driving the development of this highly theorized field: Shor's algorithm and Grover's algorithm.[91]

Peter Shor of Bell Labs designed the first quantum algorithm in 1994. Shor's algorithm allows for rapid factoring of very large numbers into their prime factors. For example, scientific estimates state that it would take a modern computer 10^{24} years to

factor a 1,000-digit number; it would take a quantum computer about 20 minutes.[92] The implications of this quantum algorithm on classic algorithms that depend on the difficulty of factoring for security, such as the widely used RSA algorithm, are immense. "The ability to break the RSA coding system will render almost all current channels of communications unsecure."[93]

Lov Grover, also of Bell Labs, invented the second quantum algorithm in 1996. Grover's algorithm allows a quantum computer to search databases of all kinds much more quickly than any capability existing today. Grover notes that the greatest benefit is gained when his algorithm is used on an unsorted database.[94] On average, it takes a classical computer $n/2$ number of searches to find a specific entry in a database of n entries. Grover's algorithm allows the same search to be done in the square root of n number of searches. For example, in a database of 1 million entries, it would take a computer today on average of 500,000 searches to find the right answer; it would take a quantum computer using Grover's algorithm only 1,000 searches. This could have implications for symmetric key algorithms such as DES, because this algorithm would allow an exhaustive search of all possible keys to occur quite rapidly.[95]

7.6.4.4 Current Status. Encouraged by the repercussions of quantum computing and the related algorithms on the security of information and cryptography, governments around the world are funding efforts to build a practical quantum computing system. The United States has many initiatives on-going. In 2001, the Defense Advanced Research Projects Agency (DARPA) of the Department of Defense launched a $100 million effort that would last five years. In addition, the National Science Foundation has $8 million in grant money for researching quantum capabilities. DARPA's Quantum Information Science and Technology initiative will now exist indefinitely; it became a fully funded and permanent program in 2006.[96] A number of other governments, primarily within Europe and Asia, are involved in quantum computation research and development. In 2000, the European Commission launched a comprehensive research effort with $20 million budgeted over three years. In Japan, the Ministry of Post and Telecommunications began an initiative in 2001 that will last 10 years with a total requested budget of $400 million. There are several commercial enterprises also involved in quantum projects. This includes IBM, Bell Labs, the Japanese firms of Fujitsu, Ltd., NEC Corporation, and Nippon Telephone and Telegraph Corporation.[97] This list is by no means exhaustive, as there are universities and other organizations throughout the world with research efforts in full swing.

Because of the worldwide effort to understand quantum computing more thoroughly, several key advancements have been made. In 1998, researchers at Los Alamos National Laboratory and MIT were able to spread a qubit over three nuclear spins of certain types of molecules. According to the experiments, spreading the information (qubit) out made it more difficult to corrupt, or decohere. The researchers were able to accomplish this using a technique called nuclear magnetic resonance (NMR), which allows the manipulation and control of a nucleus's spin. This technique allowed the researchers to use the property of entanglement to analyze indirectly the quantum information.[98]

In 2000, researchers at IBM developed a five-qubit computer, also using the nuclei of a liquid. The nuclei were programmed by radio frequency pulses and then detected by NMR techniques. Using this technique, the team was able to find the period of a particular function, or the length of the shortest interval over which it repeats its values. This problem would take a classical computer several repeated cycles to compute; the team at IBM was able to do it in one step. In 2001, a combined group of scientists from

IBM and Stanford University demonstrated Shor's algorithm and were able to find the prime factors of 15. The seven-qubit computer correctly deduced that the prime factors were 3 and 5.[99]

In February 2007, a Canadian company called D-Wave claimed to demonstrate the first commercial quantum computer. It is a "supercooled, superconducted niobium chip housing an array of 16 qubits."[100] D-Wave chose not to focus on cryptographic efforts when building the Orion, as the computer is called. Instead, Orion focuses its energy on solving pattern-matching problems and nondeterministic polynomial problems (NP-complete problems). NP-complete problems are decision problems that contain searching and optimization problems, and are used when someone needs to know if a certain solution for a certain problem exists. Examples of such problems include database searches, pattern matching, identifying diseases from symptoms, and finding matches for genetic material.[101] The company's demonstrations were done via a television feed from a remote location, due to the sensitive nature of the machine and the difficulty in transporting equipment that is cooled to just above absolute zero. Despite the demonstrations and the claims of D-Wave, scientists are skeptical that Orion is actually performing quantum computations. Even the chief executive of D-Wave said that, although all evidence indicates that Orion is performing quantum computations, there is some uncertainty. Nevertheless, D-Wave announced plans to boost the Orion to 1,000 qubits by 2008.[102]

In July 2007, scientists from NIST (United States) and the Rutherford Appleton Laboratory (United Kingdom) teamed up to explore magnetic quantum effects. This team reports having chained together "100 atoms of yttrium barium nickel oxide into a quantum spin-chain that, in effect, turn[ed] the 30-nanometer long magnetic molecule into a single element."[103] This discovery is an important step toward putting qubits onto solid-state circuits. Thirty nanometers is well beyond the atomic length scale, and it is unusual to see quantum coherence beyond the atomic level. However, the team did report stable coherent states at this size, which is large enough for the lithographic techniques used to create circuit boards and conductors of classical computers.[104]

In April 2013, researchers "successfully transmitted a secure quantum code through the atmosphere from an aircraft to a ground station." The author continues,

> "This demonstrates that quantum cryptography can be implemented as an extension to existing systems," says LMU's Sebastian Nauerth. In the experiment, single photons were sent from the aircraft to the receiver on the ground. The challenge was to ensure that the photons could be precisely directed at the telescope on the ground in spite of the impact of mechanical vibrations and air turbulence. "With the aid of rapidly movable mirrors, a targeting precision of less than 3 m over a distance of 20 km was achieved," reports Florian Moll, project leader at the DLR's Institute for Communication and Navigation. With this level of accuracy, William Tell could have hit the apple on his son's head even from a distance of 500 m.
>
> With respect to the rate of signal loss and the effects of air turbulence, the conditions encountered during the experiment were comparable to those expected for transmission via satellite. The same holds for the angular velocity of the aircraft. The success of the experiment therefore represents an important step towards secure satellite-based global communication.[105]

Even with the advances just mentioned, skeptics believe that practical quantum computers that outperform classical computers are still years, or even decades, away. After conducting many hours of research on the topic of quantum computing, this author's opinion is that it is not a matter of *if* quantum computing will become a reality but a matter of *when*. That scientists have been able to demonstrate a few theoretical

quantum computations on systems comprised of only a few qubits is highly promising. Yet scientists need to overcome many obstacles. Systems containing hundreds or thousands of qubits will be needed to perform useful computations. In addition, precise controls will be required to accomplish operations while avoiding decoherence; in fact, decoherence is perhaps the biggest obstacle to the creation of a quantum system. Until scientists can reliably measure information produced by qubits at work, it is unlikely that a practical quantum system will be built in the near future.[106]

7.6.5 Snake Oil Factor. As encryption vendors and cryptographers come to grips with the implementation and extended testing of new algorithms, it is important to note these words from the AES competition requirements:

> A complete written specification of the algorithm shall be included, consisting of all necessary mathematical equations, tables, diagrams, and parameters that are needed to implement the algorithm.

In other words, there is no secret about how the AES will make things secret, just as there is no secret about how DES works. This often strikes the crypto-novice as illogical. Why not keep the algorithm secret? Surely that will make any messages encrypted with it that much harder to decrypt. Not really. Any reliance on the secrecy of the algorithm inserts a weak link in the chain of security. Encrypting data does not guarantee that it will remain confidential. The keys must be kept secret, and the identity of persons requesting authorized access must be verified to ensure they are authentic, and so on. This is true of public key encryption as well as private key encryption.

This principle is known as Kerckhoffs' Principle, based on an 1883 publication by military cryptographer Auguste Kerckhoffs:

1. The system must be practically, if not mathematically, indecipherable;
2. It must not be required to be secret, and it must be able to fall into the hands of the enemy without inconvenience;
3. Its key must be communicable and retainable without the help of written notes, and changeable or modifiable at the will of the correspondents;
4. It must be applicable to telegraphic correspondence;
5. It must be portable, and its usage and function must not require the concourse of several people;
6. Finally, it is necessary, given the circumstances that command its application, that the system be easy to use, requiring neither mental strain nor the knowledge of a long series of rules to observe.[107]

There is no benefit to be gained by relying on an algorithm that has not been subject to open review, particularly when strong, reviewed algorithms exist. Beware of encryption vendors, or producers of any security products, that claim strength based on secret algorithms. Such claims are often a case of snake oil. (For more on bogus claims for crypto products, see Curtin's "Snake Oil FAQ," included in Section 7.8, "Further Reading," at the end of this chapter.)

7.7 STEGANOGRAPHY. Instead of scrambling data through cryptography, one can also insert data covertly into other data streams. *Steganography* (literally *covered writing* in Greek) generally uses the low-order bits of a data stream—typically an

image—to convey the cleartext. In today's high-resolution representations of color images, modifying the least significant bits of a pixel makes a negligible change in color, at least to the human eye. The steganographic software can make the changes and then extract them from the modified image.[108]

Such modified images are difficult to identify, but steganography detection tools, which rely on detecting abnormal patterns in the pixels of a carrier image, do exist.[109] For example, StegoHunt™ and StegoAnalyst software from Wetstone Technologies can identify and analyze steganographically modified data; StegoBreak can extract the cleartext from the carrier file.[110]

7.8 FURTHER READING. As stated at the outset, this chapter was not designed to be an extensive treatise on cryptography or a complete guide to the implementation of encryption technology. There are many resources available to help readers deepen their understanding of this fundamental area of information security.

Books and Articles

Bishop, M. *Computer Security: Art and Science*. Upper Saddle River, NJ: Addison-Wesley/Pearson Education, 2003.

Hinsley, F. H., and A. Stripp, eds. *Codebreakers: The Inside Story of Bletchley Park*. Oxford, UK: Oxford University Press, 2001.

Cobb, C. *Cryptography for Dummies*. Hoboken, NJ: John Wiley & Sons, 2003.

Gilbert, G., Y. S. Weinstein, and M. Hamrick. *Quantum Cryptography.* World Scientific Publications, 2013.

Goldreich, O. *Foundations of Cryptography: Volume* **I***, Basic Tools*. New York, NY: Cambridge University Press, 2007.

Juels, Ari. " Encryption Basics." In H. Bidgoli, ed., *Handbook of Information Security*, Vol. 2. Hoboken, NJ: John Wiley & Sons, 2006.

Kahn, D. *The Codebreakers: The Comprehensive History of Secret Communication from Ancient Times to the Internet*, Revised Edition. New York: Scribner, 1996.

Katz, J., and Y. Lindell. *Introduction to Modern Cryptography*, Second Edition. London: Chapman & Hall/CRC, 2014.

Mao, W. *Modern Cryptography: Theory and Practice*. Upper Saddle River, NJ: Prentice-Hall, 2003.

Mel, H. X., and D. Baker. *Cryptography Decrypted*. Addison-Wesley, Upper Saddle River, NJ 2000.

Schneier, B. *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, 2nd ed. New York, NY: John Wiley & Sons, 1996.

Seberry, J., and J. Pieprzyk. *Cryptography: An Introduction to Computer Security*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

Spillman, R. J. *Classical and Contemporary Cryptology*. Upper Saddle River, NJ: Prentice-Hall, 2004.

van Tilborg, H. C. A. & S. Jojodia, eds. *Encyclopedia of Cryptography and Security*, 2nd ed. Springer, 2013.

Yan, S. Y. *Quantum Attacks on Public-Key Cryptosystems.* Springer, 2013.

Web Resources

Arsenault, A., and S. Turner. "Internet X.509 Public Key Infrastructure: PKIX Roadmap," 2000; www.ietf.org/proceedings/00jul/I-D/pkix-roadmap-05.txt

Bacard, A. "Non-Technical PGP (Pretty Good Privacy) FAQ," 2002; www.andrebacard.com/pgp.html

Beezer, R. "Cryptography Independent Study," 2002; http://buzzard.ups.edu/courses/2002spring/iscryptos2002.html

Cate, V. Vince Cate's Cryptorebel/Cypherpunk Page, www.offshore.com.ai/security/

Cryptography Research Inc. Research Links: www.cryptography.com/resources/researchlinks.html (URL inactive)

Curtin, M. "Snake-Oil FAQ/Snake Oil Warning Signs: Encryption Software to Avoid," 1998; www.interhack.net/people/cmcurtin/snake-oil-faq.html

Electronic Frontier Foundation. "Frequently Asked Questions (FAQ) About the Electronic Frontier Foundation's 'DES Cracker' Machine," 1999; http://w2.eff.org/Privacy/Crypto/Crypto_misc/DESCracker/HTML/19980716_eff_des_faq.html

Electronic Frontier Foundation RSA. "Code-Breaking Contest Again Won by Distributed.Net and Electronic Frontier Foundation (EFF). DES Challenge III Broken in Record 22 Hours," 1999; http://w2.eff.org/Privacy/Crypto/Crypto_misc/DESCracker/HTML/19990119_deschallenge3.html

Gerck, E. "Why Is Certification Harder than It Looks?" 1999; http://mcwg.org/mcg-mirror/whycert.htm

ICSA Labs' Cryptography Community. www.icsalabs.com/icsa/main.php?pid=vjgj7567 (URL inactive)

International PGP Home Page. 2002; www.pgpi.org

Kessler, G. "An Overview of Cryptography," 2004; www.garykessler.net/library/crypto.html

PGP Home. www.pgp.com/index.php (URL inactive)

RSA Security Content Library. www.rsasecurity.com/doc_library/index.asp

Schneier, B. *Crypto-Gram* newsletter archive, 1998–2008; www.schneier.com/crypto-gram-back.html

7.9 NOTES

1. David Kahn, *The Codebreakers: The Story of Secret Writing* (New York: Scribner, 1996), pp. 980–984.
2. *The American Heritage® New Dictionary of Cultural Literacy,* 3rd ed. s.v "cryptography," http://dictionary.reference.com/browse/cryptography
3. RSA Laboratories, "What Is Cryptanalysis?" www.rsa.com/rsalabs/node.asp?id=2200 (URL inactive)
4. J. Seberry and J. Pieprzyk, *Cryptography: An Introduction to Computer Security* (Englewood Cliffs, NJ: Prentice-Hall, 1989).
5. Seberry and Pieprzyk, *Cryptography*
6. Kahn, *Codebreakers,* pp. 71–72.
7. *Encyclopaedia Britannica Online Academic Edition,* s.v. "cryptology," http://search.eb.com.library.norwich.edu/eb/article-25638. This article only available within the Norwich University system, or with a paid subscription to the Encyclopedia.
8. Brigitte Collard, "La cryptographie dans l'Antiquité gréco-romaine. III. Le chiffrement par transposition," *Folia Electronica Classica* (Louvain-la-Neuve) 7 (January-June 2004): section II(2), "Définition de la scytale"; http://bcs.fltr.ucl.ac.be/FE/07/CRYPT/Crypto44-63.html#42047

9. Brad Stark, "A Closer Look at Cryptography," Bucknell University, www.facstaff.bucknell.edu/udaepp/090/w3/brads.htm (URL inactive)
10. Kahn, *Codebreakers,* pp. 83–84.
11. "Time Table/Time-Travel through Cryptography and Cryptanalysis," www. cryptool.com/menu_zeittafel.en.html
12. Oliver Pell, "Cryptology," www.ridex.co.uk/cryptology/#_Toc439908853
13. "Time Table/Time-Travel."
14. *Encyclopaedia Britannica Online Academic Edition*, s.v. "cryptology."
15. *Encyclopaedia Britannica Online Academic Edition,* s.v. "cryptology."
16. Kahn, *Codebreakers,* p. 107.
17. *Encyclopedia Britannica Online Academic Edition,* s.v. "cryptology."
18. National Security Agency, "The Rare Books Collection: Giovanni Battista Porta," www.nsa.gov/about/cryptologic_heritage/center_crypt_history/publications/rare_books.shtml#giovanni
19. D. Salomon, *Coding for Data and Computer Communications* (New York: Springer, 2005), p. 218; http://tinyurl.com/2dsmc8
20. *Encyclopedia Britannica Online Academic Edition*, s.v. "cryptology."
21. Stallings, W. *Network and Internetwork Security Principles and Practices*. Prentice-Hall, January, 1995
22. Kevin Romano, "The Stager Ciphers and the US Military's First Cryptographic System," www.gordon.army.mil/AC/Wntr02/stager.htm
23. Cypher Research Laboratories, "A Brief History of Cryptography," www.cypher.com.au/crypto_history.htm
24. National Archives, "Teaching with Documents: The Zimmermann Telegram," www.archives.gov/education/lessons/zimmermann
25. National Archives, "Teaching with Documents."
26. Jacob Mathai, "History of Cryptography and Secrecy Systems," Fordham University, www.dsm.fordham.edu/~mathai/crypto.html#ENIGMA
27. Judson Knight, "Cryptology, History," www.espionageinfo.com/Cou-De/Cryptology-History.html
28. Knight, "Cryptology, History."
29. Oli Cooper, "Cryptography," University of Bristol, www.cs.bris.ac.uk/cooper/Cryptography/crypto.html
30. Cooper, "Cryptology."
31. Kahn, *Codebreakers,*
32. Kahn, *Codebreakers*,
33. Jacob Mathai, "History of Cryptography and Secrecy Systems," Fordham University, www.dsm.fordham.edu/~mathai/crypto.html#OneTimePad
34. Stallings, "Network and Internetwork Security"
35. SCI.CRYPT FAQ §5.2, www.faqs.org/faqs/cryptography-faq/part5
36. Kahn, *Codebreakers,* p. 979.
37. Ari Juels, "Encryption Basics," in *Handbook of Information Security*, Vol. 2, ed. H. Bidgoli, (Hoboken, NJ: John Wiley & Sons, 2006), p. 980.
38. Juels, "Encryption Basics," p. 981.
39. Juels, "Encryption Basics," p. 471.

40. SCI.CRYPT FAQ, www.faqs.org/faqs/cryptography-faq/part6/
41. Bruce Schneier, *Applied Cryptography*, 2nd ed. (New York: John Wiley & Sons, 1996), p. 461.
42. Kahn, *Codebreakers,* p. 982.
43. Juels, "Encryption Basics," p. 474.
44. Juels, "Encryption Basics," p. 474.
45. Juels, "Encryption Basics," p. 474.
46. RSA Laboratories, "What Is Diffie-Hellman?" www.rsa.com/rsalabs/node.asp?id=2248
47. Charlie Kaufman, "IPsec: IKE (Internet Key Exchange)," vol. 1, *Handbook for Information Security* (Hoboken, NJ: John Wiley & Sons, 2006), p. 974.
48. Juels, "Encryption Basics," pp. 474–475.
49. SCI.CRYPT FAQ, www.faqs.org/faqs/cryptography-faq/part6/
50. SCI.CRYPT FAQ, www.faqs.org/faqs/cryptography-faq/part6/
51. SCI.CRYPT FAQ, www.faqs.org/faqs/cryptography-faq/part6/
52. SCI.CRYPT FAQ, www.faqs.org/faqs/cryptography-faq/part6/
53. SCI.CRYPT FAQ, www.faqs.org/faqs/cryptography-faq/part6/
54. For more on this evolution, see the excellent IETF document, "Internet X.509 Public Key Infrastructure: PKIX Roadmap," by A. Arsenault and S. Turner. www.ietf.org/proceedings/00jul/I-D/pkix-roadmap-05.txt
55. A. Arsenault and S. Turner, "Internet X.509 Public Key Infrastructure: PKIX Roadmap."
56. Michael Baum, *Secure Electronic Commerce* (Prentice-Hall, 1997)
57. E. Gerck, "Why Is Certification Harder than It Looks?" 1999, http://mcwg.org/mcg-mirror/whycert.htm
58. National Security Agency, "The Case for Elliptic Curve Cryptography," www.nsa.gov/business/programs/elliptic_curve.shtml
59. National Security Agency, "The Case for Elliptic Curve Cryptography."
60. Certicom, "An Elliptic Curve Cryptography (ECC) Primer," www.certicom.com/pdfs/WP-ECCprimer_login.pdf
61. National Security Agency, "The Case for Elliptic Curve Cryptography."
62. Jacob West, "The Quantum Computer," www.cs.rice.edu/~taha/teaching/05F/210/news/2005_09_16.htm
63. West, "Quantum Computer."
64. Simon Bone and Matias Castro, "A Brief History of Quantum Computing," Imperial College, London, www.doc.ic.ac.uk/~nd/surprise_97/journal/vol4/spb3
65. Quantum Information Partners LLP, "Short History of Quantum Information Processing," www.qipartners.com/publications/Short_History_of_QC.pdf. (URL inactive)
66. West, "Quantum Computer."
67. West, "Quantum Computer."
68. Bone and Castro, "Brief History."
69. Bone and Castro, "Brief History."
70. Simon Bone, "The Hitchhiker's Guide to Quantum Computing," www.doc.ic.ac.uk/~nd/surprise_97/journal/vol1/spb3/

71. Bone, "Hitchhiker's Guide."
72. West, "Quantum Computer."
73. *American Heritage ® Dictionary of the English Language, Fourth Edition*, s.v. "quanta," http://dictionary.reference.com/browse/quantum
74. Genomics and Proteomics, "Glossary: Quantum Mechanics," www.genpromag.com/Glossary.aspx?LETTER=Q (URL inactive)
75. Stephen Jenkins, "Some Basic Ideas About Quantum Mechanics," University of Exeter, newton.ex.ac.uk/research/qsystems/people/jenkins/mbody/mbody2.html
76. West, "Quantum Computer."
77. Bone and Castro, "Brief History."
78. Bonsor and Strickland, "How Quantum Computers Work," computer. www.howstuffworks.com/quantum-computer3.htm
79. Bone and Castro, "Brief History."
80. Duncan McKimm, "Quantum Entanglement," www.abc.net.au/science/features/quantum
81. West, "Quantum Computer."
82. Bone and Castro, "Brief History."
83. Bone and Castro, "Brief History."
84. Bone and Castro, "Brief History."
85. Bone and Castro, "Brief History."
86. Bonsor and Strickland, "How Quantum Computers Work."
87. SearchSMB.com, "Entanglement," http://searchcio-midmarket.techtarget.com/definition/entanglement
88. McKimm, "Quantum Entanglement."
89. SearchSMB.com, "Entanglement."
90. Bone and Castro, "Brief History."
91. Bone and Castro, "Brief History."
92. Bone and Castro, "Brief History."
93. Bone, "Hitchhiker's Guide."
94. Lov Grover, "What's a Quantum Phone Book?" www.bell-labs.com/user/feature/archives/lkgrover
95. Bone and Castro, "Brief History."
96. Confidential Source #3, "Quantum Computing," Internal Research branch Web site, August 10, 2007.
97. Confidential Source #3, "Quantum Computing."
98. West, "Quantum Computer."
99. Bonsor and Strickland, "How Quantum Computers Work."
100. R. Colin Johnson "Quantum Computer 'Orion' Debuts," EETimes, www.eetimes.com/electronics-news/4069654/Quantum-computer-Orion-debuts
101. Johnson, "Quantum Computer 'Orion' Debuts."
102. Jordon Robertson, "Scientists Dubious of Quantum Computer Claims," abcnews.go.com/Technology/wireStory?id=2875656.
103. R. Colin Johnson, "Circuit-Sized Quantum Effect Observed," EETimes, www.eetimes.com/electronics-news/4073585/Circuit-sized-quantum-effect-observed

104. Johnson, "Circuit-Sized Quantum Effect Observed."
105. ScienceDaily. "Quantum Cryptography: On Wings of Light." *ScienceDaily Science News.* 04 13, 2013. www.sciencedaily.com/releases/2013/04/130403071950.htm (accessed May 16, 2013).
106. Johnson, "Circuit-Sized Quantum Effect Observed."
107. Petitcolas, Fabien. *la cryptographie militaire.* 2012. www.petitcolas.net/fabien/kerckhoffs/#english (accessed May 21, 2013).
108. Kessler, Gary C. *Steganography: Hiding Data Within Data.* 2002. www.garykessler.net/library/steganography.html (accessed May 21, 2013).
109. National Institute of Justice. *Digital Evidence Analysis: Steganography Detection.* 11 05, 2010. www.nij.gov/topics/forensics/evidence/digital/analysis/steganography.htm (accessed 05 21, 2013).
110. Wetstone Technologies. "StegoHunt." *Wetstone.* 2013. www.wetstonetech.com/product/stegohunt/ (accessed May 21, 2013).

CHAPTER 8

USING A COMMON LANGUAGE FOR COMPUTER SECURITY INCIDENT INFORMATION

John D. Howard

8.1 INTRODUCTION 8·1

8.2 WHY A COMMON LANGUAGE IS NEEDED 8·2

8.3 DEVELOPMENT OF THE COMMON LANGUAGE 8·3

8.4 COMPUTER SECURITY INCIDENT INFORMATION TAXONOMY 8·4

8.4.1 Events 8·4

8.4.2 Attacks 8·11

8.4.3 Full Incident Information Taxonomy 8·15

8.5 ADDITIONAL INCIDENT INFORMATION TERMS 8·16

8.5.1 Success and Failure 8·17

8.5.2 Site and Site Name 8·17

8.5.3 Other Incident Terms 8·17

8.6 HOW TO USE THE COMMON LANGUAGE 8·18

8.7 NOTES 8·20

8.1 INTRODUCTION. A computer security *incident* is some set of events that involves an attack or series of attacks at one or more sites. (See Section 8.4.3 for a more formal definition of the term "incident.") Dealing with these incidents is inevitable for individuals and organizations at all levels of computer security. A major part of dealing with these incidents is recording and receiving incident information, which almost always is in the form of relatively unstructured text files. Over time, these files can end up containing a large quantity of very valuable information. Unfortunately, the unstructured form of the information often makes incident information difficult to manage and use.

This chapter presents the results of several efforts over the last few years to develop and propose a method to handle these unstructured computer security incident records. Specifically, this chapter presents a *tool* designed to help individuals and organizations record, understand, and share computer security incident information. We call the tool the *common language for computer security incident information*. This common language contains two parts:

1. A set of "high-level" incident-related terms
2. A method of classifying incident information (a taxonomy)

The two parts of the common language, the terms and the taxonomy, are closely related. The taxonomy provides a structure that shows how most common-language terms are related. The common language is intended to help investigators improve their ability to:

- Talk more understandably with others about incidents
- Gather, organize, and record incident information
- Extract data from incident information
- Summarize, share, and compare incident information
- Use incident information to evaluate and decide on proper courses of action
- Use incident information to determine effects of actions over time

This chapter begins with a brief overview of why a common language is needed, followed by a summary of how the incident common language was developed. We then present the common language in two parts: (1) incident terms and taxonomy and (2) additional incident information terms. The final section contains information about some practical ways to use the common language.

8.2 WHY A COMMON LANGUAGE IS NEEDED. When the first edition of this *Handbook* was published more than 30 years ago, computer security was a small, obscure, academic specialty. Because there were only a few people working in the field, the handling of computer security information could largely take place in an ad hoc way. In this environment, individuals and groups developed their own terms to describe computer security information. They also developed, gathered, organized, evaluated, and exchanged their computer security information in largely unique and unstructured ways. This lack of generalization has meant that computer security information has typically not been easy to compare or combine, or sometimes even to talk about in an understandable way.

Progress over the years in agreeing on a relatively standard set of terms for computer security (a common language) has had mixed results. One problem is that many terms are not yet in widespread use. Another problem is that the terms that are in widespread use often do not have standard meanings. An example of the latter is the term "computer virus." We hear the term frequently, not only in academic forums but also in the news media and popular publications. It turns out, however, that even in academic publications, "computer virus" has no accepted definition.[1] Many authors define a computer virus to be "a code fragment that copies itself into a larger program."[2] They use the term "worm" to describe an independent program that performs a similar invasive function (e.g., the Internet Worm in 1988). But other authors use the term "computer virus" to describe *both* invasive code fragments and independent programs.

Progress in developing methods to gather, organize, evaluate, and exchange computer security information also has had limited success. For example, the original records (1988–1992) of the Computer Emergency Response Team (now the CERT Coordination Center or CERT/CC) are simply a file of email and other files sent to the CERT/CC. These messages and files were archived together in chronological order, without any other organization. After 1992, the CERT/CC and other organizations developed methods to organize and disseminate their information, but the information remains difficult to combine or compare because most of it remains almost completely textual information that is uniquely structured for the CERT/CC.

Such ad hoc terms and ad hoc ways to gather, organize, evaluate, and exchange computer security information are no longer adequate. Far too many people and organizations are involved, and there is far too much information to understand and share. Today computer security is an increasingly important, relevant, and sophisticated field of study. Numerous individuals and organizations now regularly gather and disseminate computer security information. Such information ranges all the way from the security characteristics and vulnerabilities of computers and networks, to the behavior of people and systems during security incidents—far too much information for each individual and organization to have its own unique language.

One of the key elements to making systematic progress in any field of inquiry is the development of a consistent set of terms and taxonomies (principles of classification) that are used in that field.[3] This is a necessary and natural process that leads to a growing *common language,* which enables gathering, exchanging, and comparing information. In other words, the more a field of inquiry such as computer security grows, the more a common language is needed to understand and communicate with one another.

8.3 DEVELOPMENT OF THE COMMON LANGUAGE. Two of the more significant efforts in the process of developing this common language for computer security incident information were (1) a project to classify more than 4,300 Internet security incidents completed in 1997,[4] and (2) a series of workshops in 1997 and 1998 called the *Common Language Project*. Workshop participants included people primarily from the Security and Networking Research Group at the Sandia National Laboratories, Livermore, California, and from the CERT/CC at the Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania. Additional participation and review came from people in the Department of Defense (DoD) and the National Institute of Standards and Technology (NIST).

These efforts to develop the common language were *not* efforts to develop a comprehensive dictionary of terms. Instead, the participants were trying to develop both a *minimum* set of "high-level" terms to describe computer security attacks and incidents, and a structure and classification scheme for these terms (a *taxonomy*), which could be used to classify, understand, exchange, and compare computer security attack and incident information.

Participants in the workshops hoped this common language would gain wide acceptance because of its usefulness. There is already evidence that this acceptance is taking place, particularly at incident response teams and in the DoD.

In order to be complete, logical, and useful, the common language for computer security incident information was based initially and primarily on theory (i.e., it was a priori or nonempirically based).[5] Classification of actual Internet security incident information was then used to refine and expand the language. More specifically, the common language development proceeded in six stages:

1. Records at the CERT/CC for incidents reported to them from 1988 through 1995 were examined to establish a preliminary list of terms used to describe computer security incidents.
2. The terms in this list, and their definitions, were put together into a structure (a preliminary taxonomy).
3. This preliminary taxonomy was used to classify the information in the 1988 through 1995 incident records.
4. The preliminary taxonomy and classification results were published in 1997.[6]

5. A series of workshops was conducted from 1997 through 1998 (the *Common Language Project*) to make improvements to the taxonomy and to add additional terms.
6. The results of the workshops (the "common language for security incidents") were first published in 1998.

A *taxonomy* is a classification scheme (a structure) that partitions a body of knowledge and defines the relationship of the pieces.[7] Most of the terms in this common language for security incident information are arranged in such a taxonomy, as presented in the next section. *Classification* is the process of using a taxonomy for separating and ordering. As discussed earlier, classification of information using a taxonomy is necessary for computer security incident information because of the rapidly expanding amount of information and the nature of that information (primarily text). Classification using the common-language taxonomy is discussed in the final section of this chapter.

Our experience has shown that satisfactory taxonomies have classification categories with these six characteristics[8]:

1. **Mutually exclusive**. Classifying in one category excludes all others because categories do not overlap.
2. **Exhaustive**. Taken together, the categories include all possibilities.
3. **Unambiguous**. The taxonomy is clear and precise, so that classification is not uncertain, regardless of who is doing the classifying.
4. **Repeatable**. Repeated applications result in the same classification, regardless of who is doing the classifying.
5. **Accepted**. It is logical and intuitive, so that categories can become generally approved.
6. **Useful**. The taxonomy can be used to gain insight into the field of inquiry.

These characteristics were used to develop and evaluate the common-language taxonomy. A taxonomy, however, is merely an approximation of reality, and as such, even the best taxonomy will fall short in some characteristics. This may be especially true when the characteristics of the data being classified are imprecise and uncertain, as is typical for computer security incident information. Nevertheless, classification is an important, useful, and necessary prerequisite for systematic study of incidents.

8.4 COMPUTER SECURITY INCIDENT INFORMATION TAXONOMY.

We have been able to structure most of the terms in the common language for security incident information into a taxonomy. These terms and the taxonomy are presented in this section. Additional terms that describe the more general aspects of incidents are presented in Section 8.5.

8.4.1 Events.

The operation of computers and networks involves innumerable *events*. In a general sense, an event is a discrete change of state or status of a system or device.[9] From a computer security viewpoint, these changes of state result from *actions* that are directed against specific *targets*. An example is a user taking action to log in to the user's account on a computer system. In this case, the action taken by the user is to *authenticate* to the login program by claiming to have a specific identity and

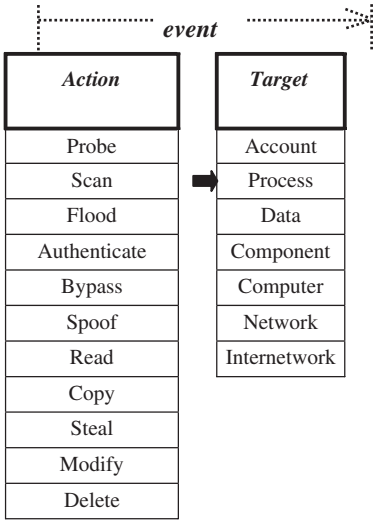


EXHIBIT 8.1 Computer and Network Events

then presenting the required verification. The target of this action would be the user's *account*. Other examples include numerous actions that can be targeted toward:

- *Data* (e.g., actions to *read*, *copy*, *modify*, *steal,* or *delete*)
- A *process* (e.g., actions to *probe*, *scan*, *authenticate*, *bypass*, or *flood* a running computer process or execution thread)
- A *component*, *computer*, *network*, or *internetwork* (e.g., actions to *scan* or *steal*)

Exhibit 8.1 presents a matrix of actions and targets that represent possible computer and network events (although not all of the possible combinations shown are feasible). A computer or network event is defined as:

> **Event**—action directed at a target that is intended to result in a change of state, or status, of the target.[10]

Several aspects of this definition are important to emphasize. First, in order for there to be an event, there must be an action that is taken, and it must be directed against a target, but the action does *not* have to succeed in actually changing the state of the target. For example, if a user enters an incorrect user name and password combination when logging in to an account, an authentication event has taken place, but the event was not successful in verifying that the user has the proper credentials to access that account.

A second important aspect is that an event represents a *practical* linkage between an action and a specific target against which the action is directed. As such, it represents the way people generally conceptualize events on computers and networks, and not all of the individual steps that actually take place during an event. For example, when a user logs in to an account, we classify the action as *authenticate* and the target as *account*. The actual action that takes place is for the user to access a process (e.g., a "login" program) in order to authenticate. Trying to depict all of the individual steps

is an unnecessary complication; the higher-level concepts presented here can describe correctly and accurately the event in a form well understood by people. In other words, it makes sense to abstract the language and its structure to the level at which people generally conceptualize the events.

By all means, supporting evidence should be presented so the evidence provides a complete idea of what happened. Stated another way, abstraction, conceptualization, and communication should be applied as close to the evidence as possible. For example, if a network switch is the target of an attack, then the target should normally be viewed as a computer or as a component (depending on the nature of the switch), and not the network, because assuming the network is the target may be an inaccurate interpretation of the evidence.

Another aspect of the definition of event is that it does not make a distinction between *authorized* and *unauthorized* actions. Most events that take place on computers or networks are both routine and authorized and, therefore, are not of concern to security professionals. Sometimes, however, an event is part of an attack or is a security concern for some other reason. This definition of event is meant to capture both authorized and unauthorized actions. For example, if a user authenticates properly, by giving the correct user identification and password combination while logging in to an account, that user is given access to that account. It may be the case, however, that this user is masquerading as the actual user, after having obtained the user identification and password from snooping on the network. Either way, this is still considered authentication.

Finally, an important aspect of events is that not all of the possible events (the action–target combinations depicted in Exhibit 8.1) are considered likely or even possible. For example, an action to *authenticate* is generally associated with an *account* or a *process* and not a different target, such as *data* or a *component*. Other examples include *read* and *copy*, which are generally targeted toward *data*; *flooding*, which is generally targeted at an *account*, *process,* or *system*; or *stealing*, which is generally targeted against *data*, a *component*, or a *computer*.

We define *action* and *target* as follows:

Action—step taken by a user or process in order to achieve a result,[11] such as to probe, scan, flood, authenticate, bypass, spoof, read, copy, steal, modify, or delete.

Target—computer or network logical entity (account, process, or data) or a physical entity (component, computer, network or internetwork).

8.4.1.1 Actions. The actions depicted in Exhibit 8.1 represent a spectrum of activities that can take place on computers and networks. An action is a step taken by a *user* or a *process* in order to achieve a result. Actions are initiated by accessing a target, where *access* is defined as:

Access—establish logical or physical communication or contact.[12]

Two actions are used to gather information about targets: *probe* and *scan*. A probe is an action to determine one or more characteristics of a specific target. This is unlike a scan, which is an action where a user or process accesses a set of targets systematically, in order to determine which targets have one or more characteristics.

"Probe" and "scan" are terms commonly used by incident response teams. As a result, they have common, accepted definitions. Despite this, there is a logical ambiguity: A scan could be viewed as multiple probes. In other words, if an attacker is testing

	Test for One or More Characteristics
Test a Single Host	Probe
Nonsystematically Test Multiple Hosts	Multiple probes
Systematically Test a Set of Hosts	Scan

EXHIBIT 8.2 Probe Compared to Scan

for one or more characteristics on multiple hosts, this can be (a) multiple attacks (all *probes*), or (b) one attack (a *scan*). This point was discussed extensively in the Common Language Project workshops, and the conclusion was that the terms in the common language should match, as much as possible, their common usage. This common usage is illustrated in Exhibit 8.2.

With probes and scans, it is usually obvious what is taking place. The attacker is either "hammering away" at one host (a *probe*), randomly testing many hosts (multiple *probes*), or using some "automatic" software to look for the same characteristic(s) systematically across a group of hosts (a *scan*). As a practical matter, incident response teams do not usually have a problem deciding what type of action they are dealing with.

One additional point about *scan* is that the term "systematic" is not meant to specify some specific pattern. The most sophisticated attackers try to disguise the systematic nature of a scan. A scan may, at first, appear to be multiple probes. For example, an attacker may randomize a scan with respect to hosts and with respect to the characteristic(s) being tested. If the attack can be determined to involve testing of one or more characteristics on a group of hosts with some common property (e.g., an Internet Protocol [IP] address range) or if tests on multiple hosts appear to be otherwise related (e.g., having a common origin in location and time), then the multiple probes should be classified as a scan.

Unlike probe or scan, an action taken to *flood* a target is not used to gather information about the target. Instead, the desired result of a flood is to overwhelm or overload the target's capacity by accessing the target repeatedly. An example is repeated requests to open connections to a port over a network or repeated requests to initiate processes on a computer. Another example is a high volume of email messages, which may exceed the resources available for the targeted account.

Authenticate is an action taken by a user to assume an identity. Authentication starts with a user accessing an authentication process, such as a login program. The user must claim to have a certain identity, such as by entering a user name. Usually verification is also required as a second authentication step. For verification, the user must prove knowledge of some secret (e.g., a password), prove the possession of some token (e.g., a secure identification card), and/or prove to have a certain characteristic (e.g., a retinal scan pattern). Authentication can be used not only to log in to an account but also to access other objects, such as to operate a process or to access a file. In other words, the target of an authentication action is the entity (e.g., account, process, or data) that the user is trying to access, not the authentication process itself.

Two general methods might be used to defeat an authentication process. First, a user could obtain a valid identification and verification pair that could be used to authenticate, even though it does not belong to that user. For example, during an

incident, an attacker might use a process operating on an Internet host computer that captures user name, password, and IP address combinations that are sent in clear text across the Internet. The attacker could then use this captured information to authenticate (log in) to accounts that belong to other users. It is important to note, as mentioned earlier, that this action is still considered *authenticate*, because the attacker presents valid identification and verification pairs, even though they have been stolen.

The second method that might be used to defeat an authentication process is to exploit a vulnerability in order to bypass the authentication process and access the target. *Bypass* is an action taken to avoid a process by using an alternative method to access a target. For example, some operating systems have vulnerabilities that an attacker could exploit to gain privileges without actually logging in to a privileged account.

As was discussed with respect to *authenticate*, an action to *bypass* does not necessarily indicate that the action is unauthorized. For example, some programmers find it useful to have a shortcut ("back-door") method to enter an account or run a process, particularly during development. In such a situation, an action to bypass may be considered authorized.

Authenticate and *bypass* are actions associated with users identifying themselves. In network communications, processes also identify themselves to each other. For example, each packet of information traveling on a network contains addresses identifying both the source and the destination, as well as other information. "Correct" information in these communications is assumed, since it is automatically generated. Thus, no action is included on the list to describe this normal situation. Incorrect information could, however, be entered into these communications. Supplying such false information is commonly called an action to *spoof*. Examples include IP spoofing, mail spoofing, and Domain Name System (DNS) spoofing.

> *Spoofing* is an active security attack in which one machine on the network masquerades as a different machine. . . . [It] disrupts the normal flow of data and may involve injecting data into the communications link between other machines. This masquerade aims to fool other machines on the network into accepting the imposter as an original, either to lure the other machines into sending it data or to allow it to alter data.[13]

Some actions are closely associated with data found on computers or networks, particularly with files: *read*, *copy*, *modify*, *steal*, and *delete*. There has been some confusion over these terms because their common usage in describing the physical world sometimes differs from their common usage describing the electronic world. For example, if I say that an attacker *stole* a computer, then you can assume I mean the attacker took possession of the target (computer) and did not leave an identical computer in that location. If I say, however, that the attacker stole a computer *file*, what does that actually mean? It is often taken to mean that the attacker *duplicated* the file and now has a copy, but also it means that the original file is still in its original location. In other words, "steal" sometimes means something different in the physical world than it does in the electronic world.

It is confusing for there to be differences in the meaning of actions in the physical world and the electronic world. Workshop participants attempted to reconcile these differences by carefully defining each term (*read*, *copy*, *modify*, *steal*, or *delete*) so it would have a very specific and mutually exclusive meaning that matches the physical-world meaning as much as possible.

Read is defined as an action to obtain the content of the data contained within a file or other data medium. This action is distinguished conceptually from the actual physical steps that may be required to read. For example, in the process of reading a computer

file, the file may be copied from a storage location into the computer's main memory and then displayed on a monitor to be read by a user. These physical steps (copy the file into memory and then onto the monitor) are not part of the abstract concept of *read*. In other words, to read a target (obtain the content in it), copying of the file is not necessarily required, and it is conceptually not included in our definition of *read*.

The same separation of concepts is included in the definition of the term "copy." In this case, we are referring to acquiring a copy of a target without deleting the original. The term "copy" does not imply that the *content* in the target is obtained, just that a *copy* has been made and obtained. To get the content, the file must be *read*. An example is copying a file from a hard disk to a floppy disk. This copying is done by duplicating the original file while leaving the original file intact. A user would have to open the file and look at the content in order to *read* it.

Copy and *read* are both different concepts from *steal*, which is an action that results in the attacker taking possession of the target and the target also becoming unavailable to the original owner or user. This definition agrees with our concepts about physical property, specifically that there is only one object that cannot be copied. For example, if someone steals a car, then that person has deprived the owner of his or her possession. When dealing with property that is in electronic form, such as a computer file, often the term "steal" is used, when *copy* is what actually is meant. The term "steal" specifically means that the original owner or user has been denied access or use of the target. On the other hand, stealing also could mean *physically* taking a floppy disk that has the file located on it or stealing an entire computer.

Two other actions involve changing the target in some way. The first are actions to *modify* a target. Examples include changing the content of a file, changing the password of an account, sending commands to change the characteristics of an operating process, or adding components to an existing system. If the target is eliminated entirely, the term "delete" is used to describe the action.

As stated earlier, differences in usage of terms between the physical world and the electronic world are undesirable. As such, we tried to be specific and consistent in our usage. The resulting set of terms is exhaustive and mutually exclusive, but goes against the grain in some common usage for the electronic world, particularly with respect to the term "steal." The situation seems unavoidable. Here are some examples that might clarify the terms:

- A user clicks on a link with the browser and sees the content of a Web page on the computer screen. We would classify this as a *read*. While what actually happens is that the content of the page is stored in volatile memory, copied to the cache on the hard drive, and displayed on the screen, from a *logical* (i.e., user) point of view, the Web page has *not* been copied (nor stolen). Now, if a user copies the content of the Web page to a file or prints it out, then the user *has* copied the Web page. Again, this would be a logical classification of the action, from the user's point of view.
- A user duplicates a file that is encrypted. We would classify this as *copy*, not *read*. In this case, the file was reproduced, but the content not obtained, so it was not read.
- A user deletes several entries in a password or group file. Should this action be described as several *delete* actions or as one action to *modify*? We would describe this action as *modify*, and the target is *data*. There is no ambiguity here because of the definition of "data." Data are defined to be either a stationary file or a file in transit (see the next section). If a user deletes a line out of the password file, then the file has been modified. The action would be described as *delete* only if

the whole file was deleted. If we had defined data to include part of a file, then we would indeed have an ambiguity.

- A user copies a file and deletes the original. We would classify this as *steal*. Although the steps actually include a *copy* followed by a *delete*, that is the electronic way of stealing a file, and therefore it is more descriptive to describe the action as *steal*.

In reality, the term "steal" is rarely used (correctly) because attackers who copy files usually do not delete the originals. The term "steal" often is used *incorrectly*, as in "stealing the source code," when in fact the correct term is *copy*.

The list of actions was hashed over in numerous group discussions, off and on, for several years before being put into the common language. Most people who participated in these discussions were not entirely happy with the list, but it is the best we have seen so far. Specifically, the list seems to capture all of the common terms with their common usage (*probe*, *scan*, *flood*, *spoof*, *copy*, *modify*, and *delete*) and the other terms are logical (to the people who participated in the discussion groups) and are necessary to make the action category exhaustive (*authenticate*, *bypass*, *read*, and *steal*).

Here is a summary of our definitions of the actions shown in Exhibit 8.1.

Probe—access a target in order to determine one or more of its characteristics.

Scan—access a set of targets systematically in order to identify which targets have one or more specific characteristics.[14]

Flood—access a target repeatedly in order to overload the target's capacity.

Authenticate—present an identity to a process and, if required, verify that identity, in order to access a target.[15]

Bypass—avoid a process by using an alternative method to access a target.[16]

Spoof—masquerade by assuming the appearance of a different entity in network communications.[17]

Read—obtain the content of data in a storage device or other data medium.[18]

Copy—reproduce a target leaving the original target unchanged.[19]

Steal—take possession of a target without leaving a copy in the original location.

Modify—change the content or characteristics of a target.[20]

Delete—remove a target or render it irretrievable.[21]

8.4.1.2 Targets. Actions are considered to be directed toward seven categories of targets. The first three of these are "logical" entities (*account*, *process,* and *data*), and the other four are "physical" entities (*component*, *computer*, *network*, and *internetwork*).

In a multiuser environment, an *account* is the domain of an individual user. This domain includes the files and processes the user is authorized to access and use. A special program that records the user's account name, password, and use restrictions controls access to the user's account. Some accounts have increased or special permissions that allow access to system accounts, other user accounts, or system files and processes, and often are called *privileged*, *superuser*, *administrator*, or *root* accounts.

Sometimes an action may be directed toward a *process*, which is a program executing on a computer or network. In addition to the program itself, the process includes the program's data and stack; its program counter, stack pointer, and other registers; and all other information needed to execute the program.[22] The action may then be to supply information to the process or command the process in some manner.

The target of an action may be *data* that are found on a computer or network. Data are representations of facts, concepts, or instructions in forms that are suitable for use by either users or processes. Data may be found in two forms: files or data in transit. *Files* are data that are designated by name and considered as a unit by the user or by a process. Commonly we think of files as being located on a storage medium, such as a storage disk, but files also may be located in the volatile or nonvolatile memory of a computer. *Data in transit* are data being transmitted across a network or otherwise emanating from some source. Examples of the latter include data transmitted between devices in a computer and data found in the electromagnetic fields that surround computer monitors, storage devices, processors, network transmission media, and the like.

Sometimes we conceptualize the target of an action as not being a *logical* entity (account, process, or data) but rather as a *physical* entity. The smallest of the physical entities is a *component*, which is one of the parts that make up a computer or network. A *network* is an interconnected or interrelated group of computers, along with the appropriate switching elements and interconnecting branches.[23] When a computer is attached to a network, it is sometimes referred to as a *host computer*. If networks are connected to each other, then they are sometimes referred to as an *internetwork*.

Here is a summary of our definitions of the targets shown in Exhibit 8.1.

Account—domain of user access on a computer or network that is controlled according to a record of information which contains the user's account name, password, and use restrictions.

Process—program in execution, consisting of the executable program, the program's data and stack, its program counter, stack pointer and other registers, and all other information needed to execute the program.[24]

Data—representations of facts, concepts, or instructions in a manner suitable for communication, interpretation, or processing by humans or by automatic means.[25] Data can be in the form of *files* in a computer's volatile memory or nonvolatile memory, or in a data storage device, or in the form of *data in transit* across a transmission medium.

Component—one of the parts that make up a computer or network.[26]

Computer—device that consists of one or more associated components, including processing units and peripheral units, that is controlled by internally stored programs and that can perform substantial computations, including numerous arithmetic operations or logic operations, without human intervention during execution. Note: may be stand-alone or may consist of several interconnected units.[27]

Network—interconnected or interrelated group of host computers, switching elements, and interconnecting branches.[28]

Internetwork—network of networks.

8.4.2 Attacks. Sometimes an event that occurs on a computer or network is part of a series of steps intended to result in something that is not authorized to happen. This event is then considered part of an *attack*. An attack has three elements.

1. It is made up a series of steps taken by an *attacker*. Among these steps is an action directed at a target (an *event*, as described in the previous section) as well as the use of some *tool* to exploit a *vulnerability*.

2. An attack is intended to achieve an *unauthorized result* as viewed from the perspective of the owner or administrator of the system involved.
3. An attack is a series of *intentional* steps initiated by the attacker. This differentiates an attack from something that is inadvertent.

We define an attack in this way:

Attack—a series of steps taken by an attacker to achieve an unauthorized result.

Exhibit 8.3 presents a matrix of possible attacks, based on our experience. Attacks have five parts that depict the logical steps an attacker must take. An attacker uses a (1) *tool* to exploit a (2) *vulnerability* to perform an (3) *action* on a (4) *target* in order to achieve an (5) *unauthorized result*. To be successful, an attacker must find one or more paths that can be connected (attacks), perhaps simultaneously or repeatedly. The first two steps in an attack, *tool* and *vulnerability*, are used to cause an *event* (*action* directed at a *target*) on a computer or network. The logical end of a successful attack is an *unauthorized result*. If the logical end of the previous steps is an *authorized* result, then an attack has not taken place.

The concept of *authorized* versus *unauthorized* is key to understanding what differentiates an attack from the normal events that occur. It is also a system-dependent concept in that what may be authorized on one system may be unauthorized on another.

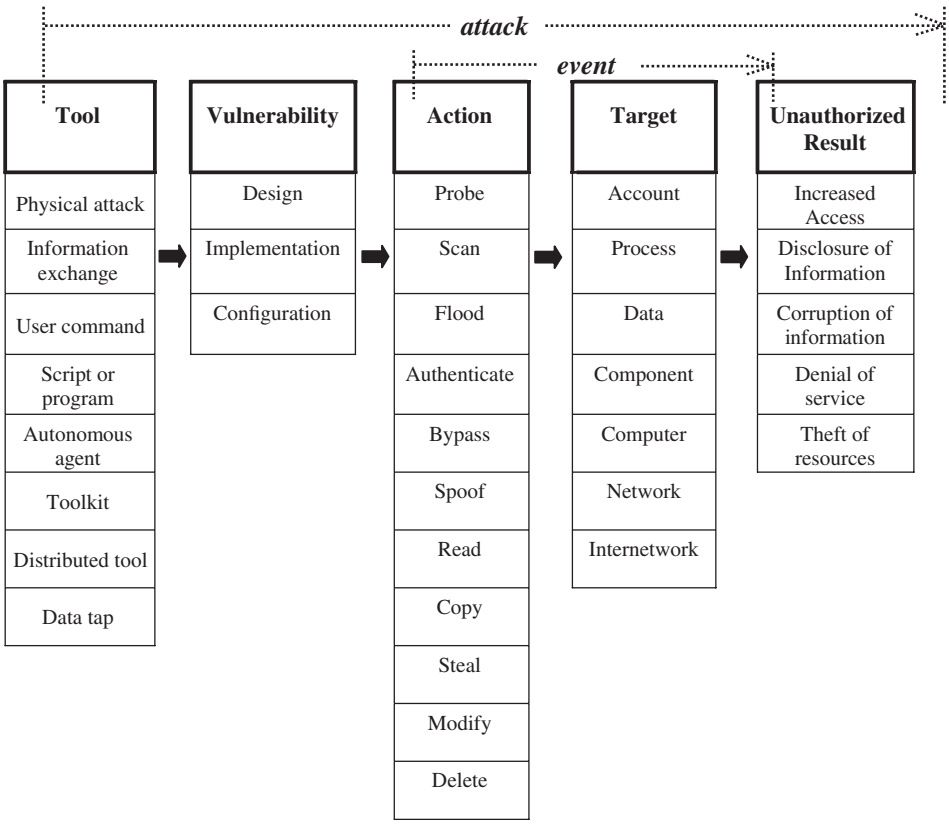


EXHIBIT 8.3 Computer and Network Attacks

For example, some services, such as anonymous File Transfer Protocol (FTP), may be enabled on some systems and not on others. Even actions that are normally viewed as hostile, such as attempts to bypass access controls to gain entry into a privileged account, may be authorized in special circumstances, such as during an approved test of system security or in the use of a "back door" during development. System owners or their administrators make the determination of what actions they consider authorized for their systems by establishing a security policy.[29] Here are the definitions for authorized and unauthorized.

Authorized—approved by the owner or administrator.

Unauthorized—not approved by the owner or administrator.

The steps *action* and *target* in Exhibit 8.1 are the two parts of an event as discussed in Section 8.4.1. The following sections discuss the other steps: *tool*, *vulnerability,* and *unauthorized result*.

8.4.2.1 Tool. The first step in the sequence that leads attackers to their unauthorized results is the *tool* used in the attack. A tool is some means that can be used to exploit a vulnerability in a computer or network. Sometimes a tool is simple, such as a user command or a physical attack. Other tools can be very sophisticated and elaborate, such as a Trojan horse program, computer virus, or distributed tool. We define *tool* in this way.

Tool—means of exploiting a computer or network vulnerability.

The term "tool" is difficult to define more specifically because of the wide variety of methods available to exploit vulnerabilities in computers and networks. When authors make lists of methods of attack, often they are actually making lists of tools. Based on our experience, these categories of tools are currently an exhaustive list. (See Exhibit 8.3)

Physical attack—means of physically stealing or damaging a computer, network, its components, or its supporting systems (e.g., air conditioning, electric power, etc.).

Information exchange—means of obtaining information either from other attackers (e.g., through an electronic bulletin board) or from the people being attacked (commonly called social engineering).

User command—means of exploiting a vulnerability by entering commands to a process through direct user input at the process interface. An example is entering UNIX commands through a telnet connection or commands at a protocol's port.

Script or program—means of exploiting a vulnerability by entering commands to a process through the execution of a file of commands (script) or a program at the process interface. Examples are a shell script to exploit a software bug, a Trojan horse log-in program, or a password-cracking program.

Autonomous agent—means of exploiting a vulnerability by using a program or program fragment that operates independently from the user. Examples are computer viruses or worms.

Toolkit—software package that contains scripts, programs, or autonomous agents that exploit vulnerabilities. An example is the widely available toolkit called *rootkit*.

Distributed tool—tool that can be distributed to multiple hosts, which then can be coordinated to anonymously perform an attack on the target host simultaneously after some time delay.

Data tap—means of monitoring the electromagnetic radiation emanating from a computer or network using an external device.

With the exception of the physical attack, information exchange, and data tap categories, each of the tool categories may contain the other tool categories *within* itself. For example, toolkits contain scripts, programs, and sometimes autonomous agents. So when a *toolkit* is used, the *script or program* category is also included. *User commands* also must be used for the initiation of scripts, programs, autonomous agents, toolkits, and distributed tools. In other words, there is an order to some of the categories in the tools block, from the simple user command category to the more sophisticated distributed tools category. In describing or classifying an attack, generally a choice must be made among several alternatives within the tools block. We chose to classify according to the *highest* category of tool used, which makes the categories mutually exclusive in practice.

8.4.2.2 Vulnerability.

To reach the desired result, an attacker must take advantage of a computer or network *vulnerability*.

Vulnerability—weakness in a system allowing unauthorized action.[30]

A vulnerability in software is an error that arises in different stages of development or use.[31] This definition can be used to give us three categories of vulnerabilities:

Design vulnerability—vulnerability inherent in the design or specification of hardware or software whereby even a perfect implementation will result in a vulnerability.

Implementation vulnerability—vulnerability resulting from an error made in the software or hardware implementation of a satisfactory design.

Configuration vulnerability—vulnerability resulting from an error in the configuration of a system, such as having system accounts with default passwords, having "world write" permission for new files, or having vulnerable services enabled.[32]

8.4.2.3 Unauthorized Result.

As shown in Exhibit 8.3, the logical end of a successful attack is an *unauthorized result*. At this point, an attacker has used a tool to exploit a vulnerability in order to cause an event to take place.

Unauthorized result—unauthorized consequence of an event.

If successful, an attack will result in one of the following[33]:

Increased access—unauthorized increase in the domain of access on a computer or network.

Disclosure of information—dissemination of information to anyone who is not authorized to access that information.

Corruption of information—unauthorized alteration of data on a computer or network.

Denial of service—intentional degradation or blocking of computer or network resources.

Theft of resources—unauthorized use of computer or network resources.

8.4.3 Full Incident Information Taxonomy. Often attacks on computers and networks occur in a distinctive group that we would classify as being part of one *incident*. What makes these attacks a distinctive group is a combination of three factors, each of which we may only have partial information about.

1. There may be one attacker, or there may be several attackers who are related in some way.
2. The attacker(s) may use similar attacks, or they may be trying to achieve a distinct or similar objective.
3. The sites involved in the attacks and the timing of the attacks may be the same or may be related.

Here is the definition of *incident*:

Incident—group of attacks that can be distinguished from other attacks because of the distinctiveness of the attackers, attacks, objectives, sites, and timing.

The three parts of an incident are shown in simplified form in Exhibit 8.4, which shows that an attacker, or group of attackers, achieves objectives by performing attacks. An incident may comprise one single attack or multiple attacks, as illustrated by the return loop in the figure.

Exhibit 8.5 shows the full incident information taxonomy. It shows the relationship of events to attacks and attacks to incidents, and suggests that preventing attackers from achieving objectives could be accomplished by ensuring that an attacker cannot make any complete connections through the seven steps depicted. For example, investigations could be conducted of suspected terrorist *attackers*, systems could be searched periodically for attacker *tools*, system *vulnerabilities* could be patched, access controls could be strengthened to prevent *actions* by an attacker to access a *targeted* account, files could be encrypted so as not to *result* in disclosure, and a public education program could be initiated to prevent terrorists from achieving an *objective* of political gain.

8.4.3.1 Attackers and Their Objectives. *People* attack computers. They do so through a variety of methods and for a variety of objectives. What distinguishes the categories of attackers is a combination of who they are and their *objectives* (what they want to accomplish).

Attacker—individual who attempts one or more attacks in order to achieve an objective.

Objective—purpose or end goal of an incident.



EXHIBIT 8.4 Simplified Computer and Network Incident

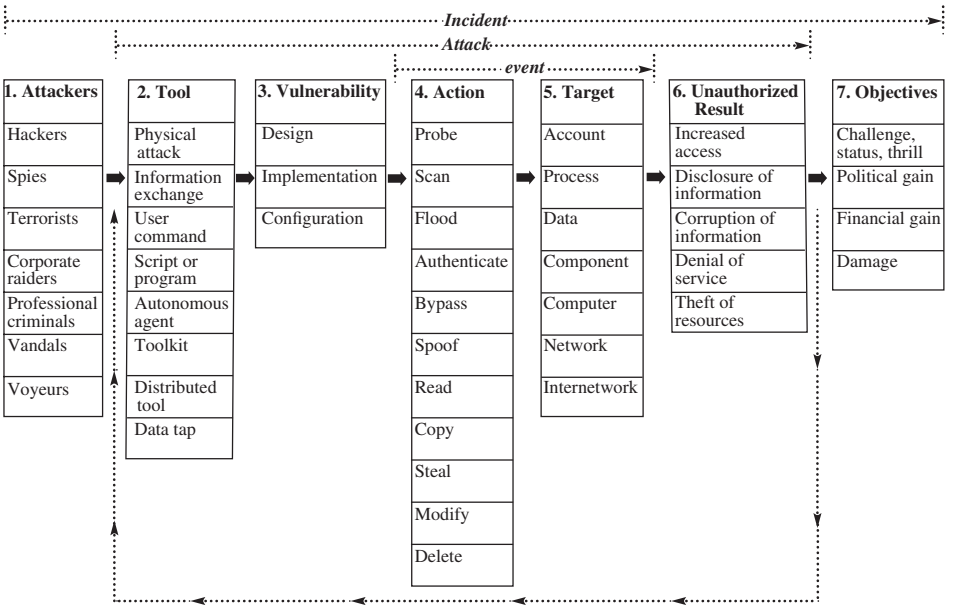


EXHIBIT 8.5 Computer and Network Incident Information Taxonomy

Based on their objectives, we have divided attackers into a number of categories:

Hackers—attackers who attack computers for challenge, status, or the thrill of obtaining access. (Note: We have elected to use the term "hacker" because it is common and widely understood. We realize that the term's more positive connotation was once more widely accepted.)

Spies—attackers who attack computers for information to be used for political gain.

Terrorists—attackers who attack computers to cause fear, for political gain.

Corporate raiders—employees (attackers) who attack competitors' computers for financial gain.

Professional criminals—attackers who attack computers for personal financial gain.

Vandals—attackers who attack computers to cause damage.

Voyeurs—attackers who attack computers for the thrill of obtaining sensitive information.

These seven categories of attackers and their four categories of objectives as shown in the leftmost and rightmost blocks of Exhibit 8.5 are fundamental to the difference between *incidents* and *attacks*. This difference is summed up in the phrase "attackers use attacks to achieve objectives."

8.5 ADDITIONAL INCIDENT INFORMATION TERMS. The taxonomy of the last section presented all of the terms in the common language for computer security that describe how attackers achieve objectives during an incident. However, some other, more general terms are required to fully describe an incident. The next sections discuss these terms.

8.5.1 Success and Failure. Information on success or failure can be recorded at several levels in the overall taxonomy. In the broadest sense, overall success or failure is an indication of whether one or more attackers have achieved one or more objectives. A narrower focus would be to determine the success or failure of an individual attack by evaluating whether the attack leads to an unauthorized result. Information on success or failure, however, may simply not be known. For example, an attempt to log in to the root or superuser account on a system may be classified as a *success* a *failure*, or as being *unknown*.

8.5.2 Site and Site Name. "Site" is the common term used to identify Internet organizations as well as physical locations. A "site" is also the organizational level of the site administrator or other authority with responsibility for the computers and networks at that location.

The term "site name" refers to a portion of the fully qualified domain name in the Internet's Domain Name System (DNS). For sites in the United States, site names generally are at the second level of the DNS tree. Examples would be *cmu.edu* or *widgets.com*. In other countries, the site name is the third or lower level of the DNS tree, such as *widgets.co.uk*. Some site names occur even farther down the DNS tree. For example, a school in Colorado might have a site name of *myschool.k12.co.us*.

Here are the definitions of site and site name.

Site—organizational level with responsibility for security events; the organizational level of the site administrator or other authority with responsibility for the computers and networks at that location.

Site name—portion of the fully qualified domain name that corresponds to a site.

Some organizations, such as larger universities and companies, are large enough to be physically divided into more than one location, with separate administration. This separation cannot easily be determined. Therefore, often these different locations must be treated as one site.

8.5.3 Other Incident Terms. Several additional terms are necessary to fully describe actual Internet incidents. The first of these terms concern dates.

Reporting date—first date that the incident was reported to a response team or other agency or individuals collecting data.

Starting date—date of the first known incident activity.

Ending date—date of the last known incident activity.

Several terms concern the sites involved.

Number of sites—overall number of sites known to have reported or otherwise to have been involved in an incident.

Reporting sites—site names of sites known to have reported an incident.

Other sites—site names of sites known to have been involved in an incident but that did not report the incident.

For most incident response teams, actual site names are considered sensitive information. In our research, in order to protect the identities of the sites associated with an incident, we sanitize the site information by coding the site names prior to public

release. An example would be to replace a site name, such as the fictitious *widgets.com*, with numbers and the upper-level domain name, such as *123.com*.

Response teams often use incident numbers to track incidents and to identify incident information.

Incident number—reference number used to track an incident or identify incident information.

The last term we found to be of use is *corrective action*, which indicates those actions taken in the aftermath of an incident. These actions could include changing passwords, reloading systems files, talking to the intruders, or even criminal prosecution. Information on corrective actions taken during or after an incident is difficult to obtain for incident response teams, since response team involvement generally is limited to the early stages of an incident. CERT/CC records indicate that the variety of corrective actions is extensive, and a taxonomy of corrective actions may be a desirable future expansion of the common language.

Corrective action—action taken during or after an incident to prevent further attacks, repair damage, or punish offenders.

8.6 HOW TO USE THE COMMON LANGUAGE. Two things are important to emphasize about using the common language for computer security incident information. First, the common language really is a high-level set of terms. As such, it will not settle all the disputes about everything discussed concerning computer security incidents. For example, the common language includes "autonomous agent" as a term (a category of tool). Autonomous agents include *computer viruses*, *worms*, and the like, regardless of how those specific terms might be defined. In other words, the common language does not try to settle disputes on what should or should not be considered a *computer virus* but rather deals at a higher level of abstraction ("autonomous agent") where, it is hoped, there can be more agreement and standardization. Stated another way, participants in the Common Language Project workshops anticipated that individuals and organizations would continue to use their own terms, which may be more specific in both meaning and use. The common language has been designed to enable these lower-level terms to be classified *within* the common language structure.

The second point to emphasize is that the common language, even though it presents a taxonomy, does not classify an incident (or individual attacks) as any *one* thing. Classifying computer security *attacks* or *incidents* is difficult because attacks and incidents are a *series of steps* that an attacker must take. In other words, attacks and incidents are not just *one* thing but rather a *series* of things. That is why I say the common language provides a taxonomy for computer security incident *information*.

An example of the problem is found in the popular and simple taxonomies often used to attempt to classify incidents. They appear as a list of single, defined terms. The following terms from Icove, Seger, and VonStorch provide an example.[34]

Covert channels	Data diddling	Degradation of service
Denial of service	Dumpster diving	Eavesdropping on emanations
Excess privileges	Harassment	IP spoofing
Logic bombs	Masquerading	Password sniffing
Salamis	Scanning	Session hijacking
Software piracy	Timing attacks	Traffic analysis
Trap doors	Trojan horses	Tunneling
Unauthorized data copying	Viruses and worms	Wiretapping

Lists of terms are *not* satisfactory taxonomies for classifying actual attacks or incidents. They fail to have most of the six characteristics of a satisfactory taxonomy. First, the terms tend not to be mutually exclusive. For example, the terms "virus" and "logic bomb" are generally found on these lists, but a virus may *contain* a logic bomb, so the categories overlap. Actual attackers generally also use multiple methods so their attacks would have to be classified into multiple categories. This makes classification ambiguous and difficult to repeat.

A more fundamental problem is that, assuming that an exhaustive and mutually exclusive list could be developed, the taxonomy would be unmanageably long and difficult to apply. It also would not indicate any relationship between different types of attacks. Finally, none of these lists has become widely accepted, partly because it is difficult to agree on the definition of terms. In fact, many different definitions of terms are in common use.

The fundamental problems with these lists (and their variations) are that most incidents involve multiple attacks, and attacks involve multiple steps. As a result, information about the typical incident must be classified in multiple categories. For example, one of the attacks in an incident might be a flood of a host resulting in a denial of service. But this same incident might involve the exploitation of a vulnerability to compromise the host computer that was the specific origin of the flood. Should this be classified as a flood? As a root compromise? As a denial-of-service attack? In reality, the incident should be classified in all of these categories. In other words, this incident has multiple classifications.

In summary, in developing the common language, we have found that, with respect to *attacks* and *incidents*, we can really *only* hope to (1) present a common set of high-level terms that are in general use and have common definitions and (2) present a logical structure to the terms that can be used to classify information *about* an incident or attack *with respect to specific categories*.

Some examples may make this clear. As discussed earlier, most of the information about actual attacks and incidents is in the form of textual records. In a typical incident record at the CERT/CC, three observations might be reported:

1. We found *rootkit* on host xxx.xxx.
2. A flood of email was sent to account xxx@xxx.xxx, which crashed the mail server.
3. We traced the attack back to a teenager in *Xyz* city, who said he was not trying to cause any damage, just trying to see if he could break in.

For observation 1, we would classify *rootkit* in the "toolkit" category under "Tool" and the hostname in the "computer" category under "Target." For observation 2, the "email flood" is a specific instantiation in the "flood" category under "Action" as well as in the "denial-of-service" category under "Unauthorized Result." There is ambiguity as to the target for observation 2: Is it the account or the computer? As a practical matter, the observations would be classified as both, since information is available on both. For observation 3, it could be inferred that this is a "hacker" seeking "challenge, status, or thrill."

What does this taxonomic process provide that is of practical value? First, the taxonomy helps us communicate to others what we have found. When we say that *rootkit* is a type of toolkit, then our common set of terms ("common language") provides us the general understanding of what we mean. When it is said that 22 percent of incidents reported to CERT/CC from 1988 through 1995 involved various problems with

passwords (a correct statistic[35]), then the taxonomy has proven useful in communicating valuable information.

The application of the taxonomy, in fact, is a four-step process that can be used to determine the biggest security problems. Specifically, the process is to:

1. Take observations from fragmentary information in incident reports.
2. Classify those observations.
3. Perform statistical studies of these data.
4. Use this information to determine the best course(s) of action.

Over time, the same process can be used to determine the effects of these actions.

Two more points are important to emphasize about this taxonomy. First, an *attack* is a process that, with enough information, is *always* classified in multiple categories. For example: in a "Tool" category, in a "Vulnerability" category, in an "Action" category, in a "Target" category, and in an "Unauthorized Result" category. Second, an *incident* can involve multiple, perhaps thousands, of attacks. As such, the information gathered in an incident theoretically could be classified correctly into *all* of the taxonomy categories.

Within these guidelines, the common language for computer security incidents has proven to be a useful and increasingly accepted tool to gather, exchange, and compare computer security information. The taxonomy itself has proven to be simple and straightforward to use.

8.7 NOTES

1. E. G. Amoroso, *Fundamentals of Computer Security Technology* (Upper Saddle River, NJ: Prentice-Hall PTR, 1994), p. 2.
2. Deborah Russell and G. T. Gangemi Sr., *Computer Security Basics* (Sebastopol, CA: O'Reilly & Associates, 1991), p. 79.
3. Bill McKelvey, *Organization Systematics: Taxonomy, Evolution, Classification* (Berkeley: University of California Press, 1982), p. 3.
4. John D. Howard, "An Analysis of Security Incidents on the Internet, 1989–1995" (PhD diss., Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, April 1997). Also available online at www.cert.org/archive/pdf/JHThesis.pdf.
5. Ivan Victor Krsul, "Software Vulnerability Analysis" (PhD diss., Computer Sciences Department, Purdue University, Lafayette, IN, May 1998), p. 12.
6. Howard, "Analysis of Security Incidents on the Internet."
7. John Radatz, ed., *The IEEE Standard Dictionary of Electrical and Electronics Terms,* 6th ed. (New York: Institute of Electrical and Electronics Engineers, 1996), p. 1087.
8. Amoroso, *Fundamentals of Computer Security Technology*, p. 34.
9. Radatz, *IEEE Standard Dictionary*, p. 373.
10. Radatz, *IEEE Standard Dictionary*, p. 373.
11. Radatz, *IEEE Standard Dictionary*, p. 11.
12. Radatz, *IEEE Standard Dictionary*, p. 5.
13. Derek Atkins et al., *Internet Security Professional Reference* (Indianapolis: New Riders Publishing, 1996), p. 258.

14. Radatz, *IEEE Standard Dictionary,* p. 947, and K. M. Jackson and J. Hruska, eds., *Computer Security Reference Book* (Boca Raton, FL: CRC Press, 1992), p. 916.
15. Merriam-Webster, *Merriam-Webster's Collegiate Dictionary,* 10th ed. (Springfield, MA: Author, 1996), pp. 77, 575, 714, and Radatz, *IEEE Standard Dictionary*, p. 57.
16. Merriam-Webster's Collegiate Dictionary, p. 157.
17. Radatz, *IEEE Standard Dictionary,* p. 630, and Atkins et al., *Internet Security*, p. 258.
18. Radatz, *IEEE Standard Dictionary*, p. 877.
19. Radatz, *IEEE Standard Dictionary,* p. 224.
20. Radatz, *IEEE Standard Dictionary*, p. 661.
21. Radatz, *IEEE Standard Dictionary*, p. 268.
22. Andrew S. Tanenbaum, *Modern Operating Systems* (Englewood Cliffs, NJ: Prentice-Hall, 1992), p. 12.
23. Radatz, *IEEE Standard Dictionary,* p. 683.
24. Tanenbaum, *Modern Operating Systems*, p. 12, and Radatz, *IEEE Standard Dictionary*, p. 822.
25. Radatz, *IEEE Standard Dictionary*, p. 250.
26. Radatz, *IEEE Standard Dictionary*, p. 189.
27. Radatz, *IEEE Standard Dictionary*, p. 192.
28. Radatz, *IEEE Standard Dictionary*, p. 683.
29. Krsul, "Software Vulnerability Analysis," pp. 5–6.
30. National Research Council, *Computers at Risk: Safe Computing in the Information Age* (Washington, DC: National Academy Press, 1991), p. 301; and Amoroso, *Fundamentals of Computer Security Technology*, p. 2.
31. Krsul, *Software Vulnerability Analysis*, pp. 10–11.
32. Atkins et al., *Internet Security*, p. 196.
33. Amoroso, *Fundamentals of Computer Security Technology*, pp. 3–4, 31; Russell and Gangemi, *Computer Security Basics*, pp. 9–10; and Frederick B. Cohen, *Protection and Security on the Information Superhighway* (New York: John Wiley & Sons, 1995), pp. 55–56.
34. David Icove, Karl Seger, and William VonStorch, *Computer Crime: A Crimefighter's Handbook* (Sebastopol, CA: O'Reilly & Associates, 1995), pp. 31–52; Cohen, *Protection and Security on the Information Superhighway*, pp. 40–54 (39 terms); and Frederick B. Cohen, "Information System Attacks: A Preliminary Classification Scheme," *Computers and Security* 16, No. 1 (1997): 29–46 (96 terms).
35. Howard, "Analysis of Security Incidents on the Internet," p. 100.

CHAPTER 9

MATHEMATICAL MODELS OF COMPUTER SECURITY

Matt Bishop

9.1 WHY MODELS ARE IMPORTANT 9·1

9.2 MODELS AND SECURITY 9·3

9.2.1 Access-Control Matrix Model 9·3
9.2.2 Harrison, Ruzzo, and Ullman and Other Results 9·5
9.2.3 Typed Access-Control Model 9·6

9.3 MODELS AND CONTROLS 9·6

9.3.1 Mandatory and Discretionary Access-Control Models 9·6
9.3.2 Originator-Controlled Access-Control Model and DRM 9·6
9.3.3 Role-Based Access-Control Models and Groups 9·7
9.3.4 Summary 9·9

9.4 CLASSIC MODELS 9·9

9.4.1 Bell-LaPadula Model 9·9
9.4.2 Biba's Strict Integrity Policy Model 9·12
9.4.3 Clark-Wilson Model 9·14
9.4.4 Chinese Wall Model 9·17
9.4.5 Summary 9·18

9.5 OTHER MODELS 9·18

9.6 CONCLUSION 9·19

9.7 FURTHER READING 9·19

9.8 NOTES 9·21

9.1 WHY MODELS ARE IMPORTANT. When you drive a new car, you look for specific items that will help you control the car: the accelerator, the brake, the shift, and the steering wheel. These exist on all cars and perform the function of speeding the car up, slowing it down, and turning it left and right. This forms a model of the car. With these items properly working, you can make a convincing argument that the model correctly describes what a car must have in order to move and be steered properly.

A model in computer security serves the same purpose. It presents a general description of a computer system (or collection of systems). The model provides a definition of "protect" (e.g., "keep confidential" or "prevent unauthorized change to") and conditions under which the protection is provided. With mathematical models, the conditions can be shown to provide the stated protection. This provides a high degree of assurance that the data and programs are protected, assuming the model is implemented correctly.

This last point is critical. To return to our car analogy, notice the phrase "with these items properly working." This also means that the average driver must be able to work

them correctly. In most, if not all, cars the model is implemented in the obvious way: The accelerator pedal is to the right of the brake pedal, and speeds the car up; the brake pedal slows it down; and turning the steering wheel moves the car to the left or right, depending on the direction that the wheel is turned. The average driver is familiar with this implementation and so can use it properly. Thus, the model and the implementation together show that this particular car can be driven.

Now, suppose that the items are implemented differently. All the items are there, but the steering wheel is locked so it cannot be turned. Even though the car has all the parts that the model requires, they do not work the way the model requires them to work. The implementation is incorrect, and the argument that the model provides does not apply to this car, because the model makes assumptions—like the steering wheel turning—that are incorrect for this car. Similarly, in all the models we present in this chapter, the reader should keep in mind the assumptions that the models make. When one applies these models to existing systems, or uses them to design new systems, one must ensure that the assumptions are met in order to gain the assurance that the model provides.

This chapter presents several mathematical models, each of which serves a different purpose. We can divide these models into several types.

The first set of models is used to determine under what conditions one can prove types of systems secure. The access-control matrix model presents a general description of a computer system that this type of model uses, and it will give some results about the decidability of security in general and for particular classes of systems.

The second type of model describes how the computer system applies controls. The mandatory access-control model and the discretionary access-control model form the basis for components of the models that follow. The originator-controlled access-control model ties control of data to the originator rather than the owner, and has obvious applications for digital rights management systems. The role-based access-control model uses job function, rather than identity, to provide controls and so can implement the principle of least privilege more effectively than many models.

The next few models describe confidentiality and integrity. The Bell-LaPadula model describes a class of systems designed to protect confidentiality and was one of the earliest, and most influential, models in computer security. The Biba model's strict integrity policy is closely related to the Bell-LaPadula model and is in widespread use today; it is applied to programs to determine when their output can be trusted. The Clark-Wilson model is also an integrity model, but it differs fundamentally from Biba's model because the Clark-Wilson model describes integrity in terms of processes and process management rather than in terms of attributes of the data.

The fourth type of model is the hybrid model. The Chinese Wall model examines conflicts of interest, and is an interesting mix of both confidentiality and integrity requirements. This type of model arises when many real-world problems are abstracted into mathematical representations, for example, when analyzing protections required for medical records and for the process of recordation of real estate.[1]

The main goal of this chapter is to provide the reader with an understanding of several of the main models in computer security, of what these models mean, and of when they are appropriate to use. An ancillary goal is to make the reader sensitive to how important assumptions in computer security are. Dorothy Denning said it clearly and succinctly in her speech when accepting the National Computer Systems Security Award in 1999:

> The lesson I learned was that security models and formal methods do not establish security. They only establish security with respect to a model, which by its very nature is extremely

simplistic compared to the system that is to be deployed, and is based on assumptions that represent current thinking. Even if a system is secure according to a model, the most common (and successful) attacks are the ones that violate the model's assumptions. Over the years, I have seen system after system defeated by people who thought of something new.[2]

Given this, the obvious question is: Why are models important? Models provide a framework for analyzing systems and for understanding where to focus our security efforts: on either validating the assumptions or ensuring that the assumptions are met in the environment in which the system exists. The mechanisms that do this may be technical; they may be procedural. Their quality determines the security of the system. So the model provides a basis for asserting that, if the mechanisms work correctly, then the system is secure—and that is far better than simply implementing security mechanisms without understanding how they work together to meet security requirements.

9.2 MODELS AND SECURITY.

Some terms recur throughout our discussion of models.

- A *subject* is an active entity, such as a process or a user.
- An *object* is a passive entity, such as a file.
- A *right* describes what a subject is allowed to do to an object; for example, the read right gives permission for a subject to read a file.
- The *protection state* of a system simply refers to the rights held by all subjects on the system.

The precise meaning of each right varies from actual system to system. For example, on Linux systems, if a process has *write* permission for a file, that process can alter the contents of the file. But if a process has *write* permission for a directory, that process can create, delete, or rename files in that directory. Similarly, having *read* rights over a process may mean the possessor can participate as a recipient of interprocess communications messages originating from that process. The point is that the meaning of the rights depends on the interpretation of the system involved. The assignment of meaning to the rights used in a mathematical model is called *instantiating the model*.

The first model we explore is the foundation for much work on the fundamental difficulty of analyzing systems to determine whether they are secure.

9.2.1 Access-Control Matrix Model.

The *access-control matrix model*[3] is perhaps the simplest model in computer security. It consists of a matrix, the rows of which correspond to subjects and the columns of which correspond to entities (subjects and objects). Each entry in the matrix contains the set of rights that the subject (row) has over the entity (column). For example, the access-control matrix in Exhibit 9.1 shows a system with two processes and two files. The first process has own rights over itself; read rights over the second process; read and execute rights over the first file; and read, write, and own rights over the second file. The second process can write to the first process; owns itself; can read, write, execute, and owns the first file; and can read the second file.

The access-control matrix captures a protection state of a system. But systems evolve; their protection state does not remain constant. So the contents of the

	Process 1	Process 2	File 1	File 2
Process 1	own	read	read, execute	read, write, own
Process 2	write	own	read, write, execute, own	read

EXHIBIT 9.1 Example Access-Control Matrix with Two Processes and Two Files

access-control matrix must change to reflect this evolution. Perhaps the simplest set of rules for changing the access-control matrix are these *primitive operations*[4]:

- **Create subject** *s* creates a new row and column, both labeled *s*
- **Create object** *o* creates a new column labeled *o*
- **Enter** *r* **into** *A*[*s*, *o*] adds the right *r* into the entry in row *s* and column *o*; it corresponds to giving the subject *s* the right *r* over the entity *o*
- **Delete** *r* **from** *A*[*s*, *o*] removes the right *r* from the entry in row *s* and column *o*; it corresponds to deleting the subject *s*'s right *r* over the entity *o*
- **Destroy subject** *s* removes the row and column labeled *s*
- **Destroy object** *o* removes the column labeled *o*

These operations can be combined into *commands*. The next command creates a file *f* and gives the process *p* read and own rights over that file:

```
command createread(p, f)
       create object f
       enter read into A[p, f]
       enter own into A[p, f]
end.
```

A *mono-operational* command consists of a single primitive operation. For example, the command

```
command grantwrite(p, f)
       enter write into A[p, f]
end.
```

which gives *p* write rights over *f,* is mono-operational.

Commands may include conditions. For example, the next command gives the subject *p* execute rights over a file *f* if *p* has read rights over *f*:

```
command grantexec(p, f)
       if read in A[p, f] then
              enter execute into A[p, f]
end.
```

If *p* does not have read rights over *f* when this command is executed, it does nothing. This command has one condition and so is called *monoconditional*. *Biconditional* commands have two conditions joined by *and*:

```
command copyread(p, q, f)
       if read in A[p, f] and own in A[p, f] then
              enter read into A[q, f]
end.
```

This command gives a subject q read rights over the object f if the subject p owns f and has read rights over f.

Commands may have conditions only at the beginning, and if the condition is false, the command terminates. Commands may contain other commands as well as primitive operations.

If all commands in a system are mono-operational, the system is said to be *mono-operational*; if all the commands are monoconditional or biconditional, then the system is said to be *monoconditional* or *biconditional*, respectively. Finally, if the system has no commands that use the *delete* or *destroy* primitive operations, the system is said to be *monotonic*.

The access-control matrix provides a theoretical basis for two widely used security mechanisms: access-control lists and capability lists. In the realm of modeling, it provides a tool to analyze the difficulty of determining how secure a system is.

9.2.2 Harrison, Ruzzo, and Ullman and Other Results.

The question of how to test whether systems are secure is critical to understanding computer security. Define *secure* in the simplest possible way: A system is secure with respect to a generic right r if that right cannot be added to an entity in the access-control matrix unless that square already contains it. In other words, a system is secure with respect to r if r cannot leak into a new entry in the access-control matrix. The question then becomes:

Safety Question. Is there an algorithm to determine whether a given system with initial state σ is secure with respect to a given right?

In the general case:

Theorem (Harrison, Ruzzo, and Ullman [HRU] Result).[5] The safety question is undecidable.

The proof is to reduce the halting problem to the safety question.[6] This means that, if the safety question were decidable, so would the halting problem be. But the undecidability of the halting problem is well known,[7] so the safety problem must also be undecidable.[8]

These results mean that one cannot develop a general algorithm for determining whether systems are secure. One can do so in limited cases, however, and the models that follow are examples of such cases. The characteristics that classes of systems must meet in order for the safety question to be decidable are not yet known fully, but for specific classes of systems, the safety question can be shown to be decidable. For example:

Theorem.[9] There is an algorithm that will determine whether mono-operational systems are secure with respect to a generic right r.

But these classes are sensitive to the commands allowed:

Theorem.[10] The safety question for monotonic systems is undecidable.

Limiting the set of commands to biconditional commands does not help:

Theorem.[11] The safety question for biconditional monotonic systems is undecidable.

But limiting them to monoconditional operations:

Theorem.[12] There is an algorithm that will determine whether monoconditional monotonic systems are secure with respect to a generic right r.

In fact, adding the *delete* primitive operation does not affect this result (although the proof is different):

Theorem.[13] There is an algorithm that will determine whether monotonic systems that do not use the *destroy* primitive operations are secure with respect to a generic right r.

9.2.3 Typed Access-Control Model. A variant of the access-control matrix model adds type to the entities. The *typed access-control matrix model*, called TAM,[14] associates a type with each entity and modifies the rules for matrix manipulation accordingly. This notion allows entities to be grouped into finer categories than merely *subject* and *object*, and enables a slightly different analysis than the HRU result suggests.

In TAM, a rule set is *acyclic* if neither an entity E nor any of its descendants can create a new entity with the same type as E. Given that definition:

Theorem.[15] There is an algorithm that will determine whether acyclic, monotonic typed matrix models are secure with respect to a generic right r.

Thus, a system being acyclic and monotonic is sufficient to make the safety question decidable. But we still do not know exactly what properties are *necessary* to make the safety question decidable.

We now turn to models that have direct application to systems and environments and that focus on more complex definitions of "secure" and the mechanisms needed to achieve them.

9.3 MODELS AND CONTROLS. Models of computer security focus on control: who can access files and resources, and what types of access are allowed. The next characterizations of these controls organize them by flexibility of use and by the roles of the entities controlling the access. These are essential to understanding how more sophisticated models work.

9.3.1 Mandatory and Discretionary Access-Control Models. Some access-control methods are rule based; that is, users have no control over them. Only the system or a special user called (for example) the *system security officer* (SSO) can change them. The government classification system works this way. Someone without a clearance is forbidden to read TOP SECRET material, even if the person who has the document wishes to allow it. This rule is called *mandatory* because it must be followed, without exception. Examples of other mandatory rules are the laws in general, which are to be followed as written, and one cannot absolve another of liability for breaking the laws; or the Multics ring-based access-control mechanism, in which accessing a data segment from below the lower bound of the segment's access bracket is forbidden regardless of the access permissions. This type of access control is called a *mandatory access control*, or MAC. These rules base the access decision on attributes of the subject and object (and possibly other information).

Other access-control methods allow the owner of the entity to control access. For example, a person who keeps a diary decides who can read it. She need not show it to anyone, and if a friend asks to read it, she can say no. Here the owner allows access to the diary at her discretion. This type of control is called *discretionary*. *Discretionary access control*, or DAC, is the most common type of access-control mechanism on computers.

Controls can be (and often are) combined. When mandatory and discretionary controls are combined to enforce a single access-control policy, the mandatory controls are applied first. If they deny access, the system denies access and the discretionary controls need never be invoked. If the mandatory rules permit access, then the discretionary controls are consulted. If both allow the accesses, access is granted.

9.3.2 Originator-Controlled Access-Control Model and DRM. Other types of access controls contain elements of both mandatory and discretionary access

controls. *Originator-controlled access control*,[16] or ORCON,[17] mechanisms allow the originator to determine who can access a resource or data.

Consider a large government research agency that produces a study of projected hoe-handle sales for the next year. The market for hoe handles is extremely volatile, and if the results of the study leak out prematurely, certain vendors will obtain a huge market advantage. But the study must be circulated to regulatory agencies so they can prepare appropriate regulations that will be in place when the study is released. Thus, the research agency must retain control of the study even as it circulates it among other groups.

More precisely, an originator-controlled access control satisfies two conditions. Suppose an object o is marked as ORCON for organization X. X decides to release o to subjects acting on behalf of another organization Y. Then

1. The subjects to whom the copies of o are given cannot release o to subjects acting on behalf of other organizations without X's consent; and
2. Any copies of o must bear these restrictions.

Consider a control that implements these requirements. In theory, mandatory access controls could solve this problem. In practice, the required rules must anticipate *all* the organizations to which the data will be made available. This requirement, combined with the need to have a separate rule for each possible set of objects and organizations that are to have access to the object, makes a mandatory access control that satisfies the requirements infeasible. But if the control were discretionary, each entity that received a copy of the study could grant access to its copy without permission of the originator; so originator-controlled access control is neither discretionary nor mandatory.

However, a combination of discretionary and mandatory access controls can implement this control. The mandatory access-control mechanisms forbid the owner from changing access permissions on an object o and require that every copy of that object have the same access-control permissions as are on o. The discretionary access control says that the originator can change the access-control permissions on any copy of o.

As an example of the use of this model in a more popular context, record companies want to control the use of their music. Conceptually, they wish to retain control over the music *after* it is sold in order to prevent owners from distributing unauthorized copies to their friends. Here the originator is the record company and the protected resource is the music.

In practice, originator-controlled access controls are difficult to implement technologically. The problem is that access-control mechanisms typically control access to *entities,* such as files, devices, and other objects. But originator-controlled access control requires that access controls be applied to *information* that is contained in the entities—a far more difficult problem for which there is not yet a generally accepted mechanism.

9.3.3 Role-Based Access-Control Models and Groups. In real life, job function often dictates access permissions. The bookkeeper of an office has free access to the company's bank accounts, whereas the sales people do not. If Anne is hired as a salesperson, she cannot access the company's funds. If she later becomes the bookkeeper, she can access those funds. So the access is conditioned not on the identity of the person but on the role that person plays.

This example illustrates *role-based access control* (RBAC).[18] It assigns a set of roles, called the *authorized roles of the subject s*, to each subject *s*. At any time, *s* may assume at most one role, called the *active role of s*. Then

Axiom. The *rule of role authorization* says that the active role of *s* must be in the set of authorized roles of *s*.

This axiom restricts *s* to assuming those roles that it is authorized to assume. Without it, *s* could assume any role, and hence do anything.

Extending this idea, let the predicate *canexec*(*s*, *c*) be true when the subject *s* can execute the command *c*.

Axiom. The *rule of role assignment* says that if *canexec*(*s*, *c*) is true for any *s* and any *c*, then *s* must have an active role.

This simply says that in order to execute a command *c*, *s* must have an active role. Without such a role, it cannot execute any commands. We also want to restrict the commands that *s* can execute; the next axiom does this.

Axiom. The *rule of transaction authorization* says that if *canexec*(*s*, *c*) is true, then only those subjects with the same role as the active role of *s* may also execute transaction.

This means that every role has a set of commands that it can execute, and if *c* is not in the set of commands that the active role of *s* can execute, then *s* cannot execute it.

As an example of the power of this model, consider two common problems: containment of roles and separation of duty. Containment of roles means that a subordinate *u* is restricted to performing a limited set of commands that a superior *s* can also perform; the superior may also perform other commands. Assign role *a* to the superior and role *b* to the subordinate; as everything a subject with active role *b* can do, a subject with active role *a* can do, we say that role *a contains* role *b*. Then we can say that if *a* is an authorized role of *s*, and *a* contains *b*, then *b* is also an authorized role of *s*. Taking this further, if a subject is authorized to assume a role that contains other (subordinate) roles, it can also assume any of the subordinate roles.

Separation of duty is a requirement that multiple entities must combine their efforts to perform a task. For example, a company may require two officers to sign a check for more than $50,000. The idea is that a single person may breach security, but two people are less likely to combine forces to breach security.[19] One way to handle separation of duty is to require that two distinct roles complete the task and make the roles mutually exclusive. More precisely, let *r* be a role and *meauth*(*r*), the mutually exclusive authorization set of *r*, be the set of roles that a subject with authorized role *r* can *never* assume. Then separation of duty is:

Axiom. The *rule of separation of duty* says that if a role *a* is in the set *meauth*(*b*), then no subject for which *a* is an authorized role may have *b* as another authorized role.

This rule is applied to a task that requires two distinct people to complete. The task is broken down into steps that two people are to complete. Each person is assigned a separate role, and each role is in the mutually exclusive authorization set of the other. This prevents either person from completing the task; they must work together, each in their respective role, to complete it.

Roles bear a resemblance to groups, but the goals of groups and roles are different. Membership in a group is defined by essentially arbitrary rules, set by the managers of the system. Membership in a role is defined by job function and is tied to a specific set of commands that are necessary to perform that job function. Thus, a role is a type of group, but a group is broader than a role and need not be tied to any particular set of commands or functions.

9.3.4 Summary. The four types of access controls discussed in this section have different focuses. Mandatory, discretionary, and originator-controlled access controls are data-centric, determining access based on the nature or attributes of the data. Role-based access control focuses on the subject's needs. The difference is fundamental.

The principle of least privilege[20] says that subjects should have no more privileges than necessary to perform their tasks. Role-based access control, if implemented properly, does this by constraining the set of commands that a subject can execute. The other three controls do this by setting attributes on the data to control access to the data rather than by restricting commands. Mandatory access controls have the attributes set by a system security officer or other trusted process; discretionary access controls, by the owner of the object; and originator-controlled access controls, by the creator or originator of the data.

As noted, these mechanisms can be combined to make the controls easier to use and more precise in application. We now discuss several models that do so.

9.4 CLASSIC MODELS.

Three models have played an important role in the development of computer security. The Bell-LaPadula model, one of the earliest formal models in computer security, influenced the development of much computer security technology, and it is still in widespread use. Biba, its analog for integrity, now plays an important role in program analysis. The Clark-Wilson model describes many commercial practices to preserve integrity of data. We examine each of these models in this section.

9.4.1 Bell-LaPadula Model. The Bell-LaPadula model[21] is a formalization of the famous government classification system using UNCLASSIFIED, CONFIDENTIAL, SECRET, and TOP SECRET levels. We begin by using those four levels to explain the ideas underlying the model and then augment those levels to present the full model. Because the model involves multiple levels, it is an example of a *multilevel security model*.

The four-level version of the model assumes that the levels are ordered from lowest to highest as UNCLASSIFIED, CONFIDENTIAL, SECRET, and TOP SECRET. Objects are assigned levels based on their sensitivity. An object at a higher level is more sensitive than an object at a lower level. Subjects are assigned levels based on what objects they can access. A subject is *cleared* into a level, and that level is called the subject's *security clearance*. An object is *classified* at a level, and that level is called the object's *security classification*. The goal of the classification system is to prevent information from leaking, or flowing downward (e.g., a subject at CONFIDENTIAL should not be able to read information classified TOP SECRET).

For convenience, we write $level(s)$ for a subject's security clearance and $level(o)$ for an object's security classification. The name of the classification is called a *label*. So an object classified at TOP SECRET has the label TOP SECRET.

Suppose Tom is cleared into the SECRET level. Three documents, called Paper, Article, and Book, are classified as CONFIDENTIAL, SECRET, and TOP SECRET, respectively. As Tom's clearance is lower than Book's classification, he cannot read Book. As his clearance is equal to or greater than Article's and Paper's classification, he can read them.

Definition. The *simple security property* says that a subject s can read an object o if and only if $level(o) \leq level(s)$.

This is sometimes called the *no-reads-up* rule, and it is a mandatory access control.

But that is insufficient to prevent information from flowing downward. Suppose Donna is cleared into the CONFIDENTIAL level. By the simple security property, she cannot read Article because

level(Article) = SECRET > CONFIDENTIAL = *level*(Donna).

But Tom can read the information in Article and write it on Paper. And Donna can read Paper. Thus, SECRET information has leaked to a subject with CONFIDENTIAL clearance.

To prevent this, Tom must be prevented from writing to Paper:

Definition. The **-property* says that a subject s can write an object o if and only if $level(s) \leq level(o)$.

This is sometimes called the *no-writes-down* rule, and it too is a mandatory access control. It is also known as the *star property* and the *confinement property.*

Under this rule, as *level*(Tom) = SECRET > *level*(Paper), Tom cannot write to Paper. This solves the problem.

Finally, the Bell-LaPadula model allows owners of objects to use discretionary access controls:

Definition. The *discretionary security property* says that a subject s can read an object o only if the access-control matrix entry for s and o contains the read right.

So, in order to determine whether Tom can read Paper, the system checks the simple security property and the discretionary security problem. As both hold for Tom and Paper, Tom can read Paper. Similarly, the system checks the *-property to determine whether Tom can write to Paper. As the *-property does not hold for Tom and Paper, Tom cannot write to Paper. Note that the discretionary security property need not be checked, because the relevant mandatory access-control property (the *-property) denies access.

The basic security theorem states that, if a system starts in a secure state, and every operation obeys the three properties, then the system remains secure:

Basic Security Theorem. Let a system Σ have a secure initial state σ_0. Further, let every command in this system obey the simple security property, the *-property, and the discretionary security property. Then every state σ_i, $i \geq 0$, is also secure.

We can generalize this to an arbitrary number of levels. Let $L_{0, \ldots}$, L_n be a set of security levels that are linearly ordered (i.e., $L_0 < \ldots < L_n$). Then the simple security property, the *-property, and the discretionary security property all apply, as does the Basic Security Theorem. This allows us to have many more than the four levels described.

Now suppose Erin works for the European Department of a government agency, and Don works for the Asia Department for the same agency. Erin and Don are both cleared for SECRET. But some information Erin will see is information that Don has no need to know, and vice versa. Introducing additional security levels will not help here, because then either Don would be able to read all of the documents that Erin could, or vice versa. We need an alternate mechanism.

The alternate mechanism is an expansion of the idea of "security level." We define a *category* to be a kind of information. A *security compartment* is a pair (*level*, *category set*) and plays the role that the security level did previously.

As an example, suppose the category for the European Department is EUR, and the category for the Asia Department is ASIA. Erin will be cleared into the compartment (SECRET, {EUR}) and Don into the compartment (SECRET, {ASIA}). Documents have security compartments as well. The paper EurDoc may be classified as

(CONFIDENTIAL, {EUR}), and the paper AsiaDoc may be (SECRET, {ASIA}). The paper EurAsiaDoc contains information about both Europe and Asia, and so would be in compartment (SECRET, {EUR, ASIA}). As before, we write *level*(Erin) = (SECRET, {EUR}), *level*(EurDoc) = (CONFIDENTIAL, {EUR}), and *level*(EurAsiaDoc) = (SECRET, {EUR, ASIA}).

Next, we must define the analog to "greater than." As noted earlier, security compartments are no longer linearly ordered, because not every pair of compartments can be compared. For example, Don's compartment is not "greater" than Erin's, and Erin's is not "greater" than Don's. But the classification of EurAsiaDoc is clearly "greater" than that of both Don and Erin.

We compare compartments using the relation *dom*, for "dominates."

Definition. Let L and L' be security levels and let C and C' be category sets. Then

$$(L,\ C)\ dom\ (L',\ C') \text{ if and only if } L' = L \text{ and } C' \subseteq C$$

The *dom* relation plays the role that "greater than or equal to" did for security levels. Continuing our example, *level*(Erin) = (SECRET, {EUR}), *dom* (CONFIDENTIAL, {EUR}) = *level*(EurDoc), and *level*(EurAsiaDoc) = (SECRET, {EUR, ASIA}) *dom* (SECRET, {EUR}) = *level*(Erin).

We now reformulate the simple security property and *-property in terms of *dom*:

Definition. The *simple security property* says that a subject s can read an object o if and only if *level*(s) *dom level*(o).

Definition. The **-property* says that a subject s can write to an object o if and only if *level*(o) *dom level*(s).

In our example, assume the discretionary access controls are set to allow any subject all types of access. In that case, as *level*(Erin) *dom level*(EurDoc), Erin can read EurDoc (by the simple security property) but not write EurDoc (by the *-property). Conversely, as *level*(EurAsiaDoc) *dom level*(Erin), Erin cannot read EurAsiaDoc (by the simple security property) but can write to EurAsiaDoc (by the *-property).

A logical question is how to determine the highest security compartment that both Erin and Don can read and the lowest that both can write. In order to do this, we must review some properties of *dom*.

First, note that *level*(s) *dom level*(s); that is, *dom* is reflexive. The relation is also antisymmetric, because if both *level*(s) *dom level*(o) and *level*(o) *dom level*(s) are true, then *level*(s) = *level*(o). It is transitive, because if *level*(s_1) *dom level*(o) and *level*(o) *dom level* (s_2), then *level*(s_1) *dom level* (s_2).

We also define the *greatest lower bound* (*glb*) of two compartments as:

Definition. Let $A = (L, C)$ and $B = (L', C')$. Then $glb(A, B) = (min(L, L'), C \cap C')$.

This answers the question of the highest security compartment that two subjects s and s' can read an object in. It is *glb*(*level*(s), *level*(s')). For example, Don and Erin can both read objects in:

$$glb(level(\text{Don}),\ level(\text{Erin})) = (\text{SECRET},\ \chi;).$$

This makes sense because Don cannot read an object in any compartment except those with the category set {ASIA} or the empty set, and Erin can only read objects in a compartment with the category set {EUR} or the empty set. Both are at the SECRET level, so the compartment must also be at the SECRET level.

We can define the *least upper bound* (*lub*) of two compartments analogously:

Definition. Let $A = (L, C)$ and $B = (L', C')$. Then $lub(A, B) = (max(L, L'), C \cup C')$.

We can now determine the lowest security compartment into which two subjects s and s' can write. It is $lub(level(s), level(s'))$. For example, Don and Erin can both write to objects in:

$$glb((level(\text{Don}),\ level(\text{Erin})) = (\text{SECRET},\ \{\text{EUR},\ \text{ASIA}\}).$$

This makes sense because Don cannot write to an object in any compartment except those with ASIA in the category set, and Erin can only write to objects in a compartment with EUR in the category set. The smallest category set meeting both these requirements is {EUR, ASIA}. Both are at the SECRET level, so the compartment must also be at the SECRET level.

The five properties of *dom* (reflexive, antisymmetric, transitive, existence of a least upper bound for every pair of elements, and existence of a greatest lower bound for every pair of elements) mean that the security compartments form a mathematical structure called a *lattice*. This has useful theoretical properties, and is important enough so models exhibiting this type of structure are called *lattice models*.

When the model is implemented on a system, the developers often make some modifications. By far the most common one is to restrict writing to the current compartment or to within a limited set of compartments. This prevents confidential information from being altered by those who cannot read it. The structure of the model can also be used to implement protections against malicious programs that alter files, such as system binaries. To prevent this, place the system binaries in a compartment that is dominated by those compartments assigned to users. By the simple security property, then users can read the system binaries, but by the *-property, users cannot write them. Hence, if a computer virus infects a user's programs or documents,[22] it can spread within that user's compartment but not to system binaries.

The Bell-LaPadula model is the basis for several other models. We explore one of its variants that models integrity rather than confidentiality.

9.4.2 Biba's Strict Integrity Policy Model. Biba's strict integrity policy model,[23] usually called *Biba's model*, is the mathematical dual of the Bell-LaPadula model.

Consider the issue of trustworthiness. When a highly trustworthy process reads data from an untrusted file and acts based on that data, the process is no longer trustworthy—as the saying goes, "garbage in, garbage out." But if a process reads data more trustworthy than the process, the trustworthiness of that process does not change. In essence, the trustworthiness of the result is as trustworthy as the least trustworthy of the process and the data.

Define a set of *integrity classes* in the same way that we defined security compartments for the Bell-LaPadula model, and let *i-level*(s) be the integrity compartment of s. Then the preceding text says that "reads down" (a trustworthy process reading untrustworthy data) should be banned, because it reduces the trustworthiness of the process. But "reads up" is allowed, because it does not affect the trustworthiness of the process. This is exactly the opposite of the simple security property.

Definition. The *simple integrity property* says that a subject s can read an object o if and only if *i-level*(o) *dom i-level*(s).

This definition captures the notion of allowing "reads up" and disallowing "reads down."

Similarly, if a trustworthy process writes data to an untrustworthy file, the trustworthiness of the file may (or may not) increase. But if an untrustworthy process writes

data to a trustworthy file, the trustworthiness of that file drops. *s* "writes down" should be allowed and "writes up" forbidden.

Definition. The **-integrity property* says that a subject *s* can write to an object *o* if and only if *i-level*(*s*) *dom i-level*(*o*).

This property blocks attempts to "write up" while allowing "writes down."

A third property relates to execution of subprocesses. Suppose process *date* wants to execute the command *time* as a subprocess. If the integrity compartment of *date* dominates that of *time*, then any information *date* passes to *time* is passed to a less trustworthy process, and hence is allowed under the *-integrity property. But if the integrity compartment of *time* dominates that of *date*, then the *-integrity property is violated. Hence:

Definition. The *execution integrity property* says that a subject *s* can execute a subject *s′* if and only if *i-level*(*s′*) *dom i-level*(*s*).

Given these three properties, one can show:

Theorem. If information can be transferred from object o_1 to object o_n, then by the simple integrity property, the *-integrity property, and the execution integrity property, *i-level*(o_1) *dom i-level*(o_n).

In other words, if all the rules of Biba's model are followed, the integrity of information cannot be corrupted because information can never flow from a less trustworthy object to a more trustworthy object.

This model suggests a method for analyzing programs to prevent security breaches. When the program runs, it reads data from a variety of sources: itself, the system, the network, and the user. Some of these sources are trustworthy, such as the process itself and the system. The user and the network are under the control of ordinary users (or remote users) and so are less trustworthy. So, apply Biba's model with two integrity compartments, (UNTAINTED, Ø) (this means the set of categories in the compartment is empty) and (TAINTED, Ø), where (UNTAINTED, Ø) *dom* (TAINTED, Ø). For notational convenience, we shall write (UNTAINTED, Ø) as UNTAINTED and (TAINTED, Ø) as TAINTED; and dom as ≥. Thus, UNTAINTED ≥ TAINTED.

The technique works with either static or dynamic analysis but is usually used for dynamic analysis. In this mode, all constants are assigned the integrity label UNTAINTED. Variables are assigned labels based on the data flows within the program. For example, in an assignment, the integrity label of the variable being assigned to is set to the integrity label of the expression assigned to it. When UNTAINTED and TAINTED variables are mixed in the expression, the integrity label of the expression is TAINTED. If a variable is assigned a value from an untrusted source, the integrity label of the variable is set to TAINTED.

When data are used as (for example) parameters of system calls or library functions, the system checks that the integrity label of the variable dominates that of the parameter. If it does not, the program takes some action, such as aborting, or logging a warning, or throwing an exception. This action either prevents an exploit or alerts the administrator of the attack.

For example, suppose a programmer wishes to prevent a format string attack. This is an attack that exploits a vulnerability in the C printing function *printf*. The first argument to *printf* is a format string, and the contents of that string determine how many other arguments *printf* expects. By manipulating the contents of a format string, an attacker can overwrite values of variables and corrupt the stack, causing the program to malfunction—usually to the attacker's benefit. The key step of the attack is to input an unexpected value for the format string. Here is a code fragment with the flaw:

```
if (fgets(buf, sizeof(buf), stdin) != NULL) printf(buf);
```

This reads a line of characters from the input into an array *buf* and immediately prints the contents of the array. If the input is "xyzzy%n", then some element of the stack will be overwritten by the value 5.[24] Hence, the first parameter to *printf* must always have integrity class UNTAINTED.

Under this analysis technique, when the input function *fgets* is executed, the variable *buf* would be assigned an integrity label of TAINTED, because the input (which is untrusted) is stored in it. Then, at the call to *printf*, the integrity class of *buf* is compared to that required for the first parameter of *printf*. The former is TAINTED; the latter is UNTAINTED. But we require that the variable's integrity class (TAINTED) dominate that of the parameter (UNTAINTED), and here TAINTED $\leq$ UNTAINTED. Hence, the analysis has found a disallowed flow and acts accordingly.

9.4.3 Clark-Wilson Model. Lipner[25] identified five requirements for commercial integrity models:

1. Users may not write their own programs to manipulate trusted data. Instead, they must use programs authorized to access that data.
2. Programmers develop and test programs on nonproduction systems, using nonproduction copies of production data if necessary.
3. Moving a program from nonproduction systems to production systems requires a special process.
4. That special process must be controlled and audited.
5. Managers and system auditors must have access to system logs and the system's current state.

Biba's model can be instantiated to meet the first and last conditions by appropriate assignment of integrity levels, but the other three focus on integrity of processes. Hence, while Biba's model works well for some problems of integrity, it does not satisfy these requirements for a commercial integrity model.

The Clark-Wilson model[26] was developed to describe processes within many commercial firms. There are several specialized terms and concepts needed to understand the Clark-Wilson mode; these are best introduced using an example:

- Consider a bank. If D are the day's deposits, W the day's withdrawals, I the amount of money in bank accounts at the beginning of the day, and F the amount of money in bank accounts at the end of the day, those values must satisfy the constraint $I + D - W = F$.
- This is called an *integrity constraint* because, if the system (the set of bank accounts) does not satisfy it, the bank's integrity has been violated.
- If the system does satisfy its integrity constraints, it is said to be in a *consistent* state.
- When in operation, the system moves from one consistent state to another. The operations that do this are called *well-formed transactions*. For example, if a customer transfers money from one account to another, the transfer is the well-formed transaction. Its component actions (withdrawal from the first account and deposit in the second) individually are not well-formed transactions, because if only one completes, the system will be in an inconsistent state.

- Procedures that verify that all integrity constraints are satisfied are called *integrity verification procedures* (IVPs).
- Data that must satisfy integrity constraints are called *constrained data items* (CDIs), and when they satisfy the constraints are said to be in a *valid* state.
- All other data are called *unconstrained data items* (UDIs).
- In addition to integrity constraints on the data, the functions implementing the well-formed transactions themselves are constrained. They must be *certified* to be well formed and to be implemented correctly. Such a function is called a *transformation procedure* (TP).

The model provides nine rules, five of which relate to the certification of data and TPs and four of which describe how the implementation of the model must enforce the certifications.

The first rule captures the requirement that the system be in a consistent state:

Certification Rule 1. An IVP must ensure that the system is in a consistent state.

The relation *certified* associates some set of CDIs with a TP that transforms those CDIs from one valid state to a (possibly different) valid state. The second rule captures this.

Certification Rule 2. For some set of associated CDIs, a TP transforms those CDIs from a valid state to a (possibly different) valid state.

The first enforcement rule ensures that the system keeps track of the *certified* relation and prevents any TP from executing with a CDI not in its associated *certified* set:

Enforcement Rule 1. The system must maintain the *certified* relation, and ensure that only TPs certified to run on a CDI manipulates that CDI.

In a typical firm, the set of users who can use a TP is restricted. For example, in a bank, a teller cannot move millions of dollars from one bank to another; doing that requires a bank officer. The second enforcement rule ensures that only authorized users can run TPs on CDIs by defining a relation *allowed* that associates a user, a TP, and the set of CDIs that the TP can access on that user's behalf:

Enforcement Rule 2. The system must associate a user with each TP and set of CDIs. The TP may access those CDIs on behalf of the associated user. If a user is not associated with a particular TP and set of CDIs, then the TP cannot access those CDIs on behalf of that user.

This implies that the system can correctly identify users. The next rule enforces this:

Enforcement Rule 3. The system must authenticate each user attempting to execute a TP.

This ensures that the identity of a person trying to execute a TP is correctly bound to the corresponding user identity within the computer. The form of authentication is left up to the instantiation of the model, because differing environments suggest different authentication requirements. For example, a bank officer may use a biometric device and a password to authenticate herself to the computer that moves millions of dollars; a teller whose actions are restricted to smaller amounts of money may need only to supply a password.

Separation of duty, already discussed, is a key consideration in many commercial operations. The Clark-Wilson model captures it in the next rule:

Certification Rule 3. The *allowed* relation must meet the requirements imposed by separation of duty.

A cardinal principle of commercial integrity is that the operations must be auditable. This requires logging of enough information to determine what the transaction did. The next rule captures this requirement:

Certification Rule 4. All TPs must append enough information to reconstruct the operation to an append-only CDI.

The append-only CDI is, of course, a log.

So far we have considered all inputs to TPs to be CDIs. Unfortunately, that is infeasible. In our bank example, the teller will enter account information and deposit and withdrawal figures; but those are not CDIs; the teller may mistype something. Before a TP can use that information, it must be vetted to ensure it will enable the TP to work correctly. The last certification rule captures this:

Certification Rule 5. A TP that takes a UDI as input must perform either a well-formed transaction or no transaction for any value of the UDI. Thus, it either rejects the UDI or transforms it into a CDI.

This also covers poorly crafted TPs; if the input can exploit vulnerabilities in the TP to cause it to act in unexpected ways, it cannot be certified under this rule.

Within the model lies a possible conflict. In the preceding rules, one user could certify a TP to operate on a CDI and then execute the TP on that CDI. The problem is that a malicious user may certify a TP that does not perform a well-formed transaction, causing the system to violate the integrity constraints. Clearly, an application of the principle of separation of duty would solve this problem, and indeed the last rule in the model does just that:

Enforcement Rule 4. Only the certifier of a TP may change the *certified* relation for that TP. Further, no certifier of a TP, or of any CDI associated with that TP, may execute the TP on the associated CDI.

This separates the ability to certify a TP from the ability to execute that TP and the ability to certify a CDI for a given TP from the ability to execute that TP on that CDI. This enforces the separation of duty requirement.

Now, revisit Lipner's requirements for commercial integrity models. The TPs correspond to Lipner's programs and the CDIs to the production data. To meet requirement 1, the Clark-Wilson certifiers need to be trusted, and ordinary users cannot certify either TPs or CDIs. Then Enforcement Rule 4 and Certification Rule 5 enforce this requirement. Requirement 2 is met by not certifying the development programs; as they are not TPs, they cannot be run on production data. The "special process" in requirement 3 is a TP. Certification Rule 4 describes a log; the special process in requirement 3 being a TP, it will append information to the log that can be audited. Further, the TP is by definition a controlled process, and Enforcement Rule 4 and Certification Rule 5 control its execution. Before the installation, the program being installed is a UDI; after it is installed, it is a CDI (and a TP). Thus, requirement 4 is satisfied. Finally, the Clark-Wilson model has a log that captures all aspects of what a TP does, and that is the log the managers and auditors will have access to. They also have access to the system state because they can run an IVP to check its integrity. Thus, Lipner's requirement 5 is met. So the Clark-Wilson model is indeed a satisfactory commercial integrity model.

This model is important for two reasons. First, it captures the way most commercial firms work, including applying separation of duty (something that Biba's model does not capture well). Second, it separates the notions of certification and enforcement. Enforcement typically can be done within the instantiation of the model. But the model cannot enforce how certification is done; it can only require that a certifier claim to have done it. This is true of all models, of course, but the Clark-Wilson model specifically states the assumptions it makes about certification.

9.4.4 Chinese Wall Model. Sometimes called the *Brewer-Nash model*, the goal of the Chinese Wall model[27] is to prevent conflicts of interest. It does so by grouping objects that belong to the same company into *company data sets* and company data sets into *conflict-of-interest classes*. If two companies (represented by their associated company data sets) are in the same conflict-of-interest class, then a lawyer or stockbroker representing both would have a conflict of interest. The rules of the model ensure that a subject can read only one company data set in each conflict-of-interest class.

In general, objects are documents or resources that contain information that the company wishes to (or is required to) keep secret. There is, however, an exception. Companies release data publicly, in the form of annual reports; that information is carefully *sanitized* to remove all confidential content. To reflect business practice, the model must allow *all* subjects to see that data. The model therefore defines a conflict-of-interest class called the *sanitized class* that has one company data set holding *only* objects containing sanitized data.

Now consider a subject reading an object. If the subject has never read any object in the object's conflict-of-interest class, reading the object presents no conflict of interest. If the subject has read an object in the same company data set, then the only information that the subject has seen in that conflict-of-interest class is from the same company as the object it is trying to read, which is allowed. But if the subject has read an object in the same conflict-of-interest class but a *different* company data set, then were the new read request granted, the subject would have read information from two different companies for which there is a conflict of interest—exactly what the model is trying to prevent. So that is disallowed.

The next rule summarizes this:

Definition. The *CW-simple security property* says that a subject *s* can read an object *o* if and only if either:

1. *s* has not read any other object in *o*'s conflict-of-interest class; or
2. The only objects in *o*'s conflict-of-interest class that *s* has read are all in *o*'s company data set.

To see why this works, suppose all banks are in the same conflict-of-interest class. A stockbroker represents The Big Bank. She is approached to represent The Bigger Bank. If she agreed, she would need access to The Bigger Bank's information, specifically the objects in The Bigger Bank's company data set. But that would mean she could read objects from two company data sets in the same conflict-of-interest class, something the CW-simple security property forbids. The temporal element of the model is important; even if she resigned her representation of The Big Bank, she cannot represent The Bigger Bank because condition 2 of the CW-simple security property considers all objects she previously read. This makes sense, because she has had access to The Big Bank, and could unintentionally compromise the interests of her previous employer while representing The Bigger Bank.

The CW-simple security property implicitly says that *s* can read any sanitized object. To see this, note that if *s* has never read a sanitized object, condition 1 holds. If *s* has read a sanitized object, then condition 2 holds because all sanitized objects are in the same company data set.

Writing poses another problem. Suppose Barbara represents The Big Bank, and Percival works for The Bigger Bank. Both also represent The Biggest Toy Company, which—not being a financial institution—is in a different conflict-of-interest class from

either bank. Thus, there is no conflict of interest in either Barbara's or Percival's representation of a bank and the toy company. But there is a path along which information can flow from Barbara to Percival, and vice versa, that enables a conflict of interest to occur. Barbara can read information from an object in The Big Bank's company data set and write it to an object in The Biggest Toy Company's company data set. Percival can read the information from the object in The Biggest Toy Company's company data set, thereby effectively giving him access to The Big Bank's information—which is a conflict of interest. That he needs Barbara's help does not detract from the problem. The goal of the model requires that this conspiracy be prevented. The next rule does so:

Definition. The *CW-*-property* says that a subject *s* can write to an object *o* if and only if both of the following conditions are met:

1. The CW-simple security property allows *s* to read *o*; and
2. All unsanitized objects that *s* can read are in the same company data set as *o*.

Now Barbara can read objects in both The Big Bank's company data set and The Biggest Toy Company's data set. But when she tries to write to The Biggest Toy Company's data set, the CW-*-property prevents her from doing so as condition 2 is not met (because she can read an object in The Big Bank's company data set).

This also accounts for sanitized objects. Suppose that Skyler represents The Biggest Toy Company and no other company. He can also read information from the sanitized class. When he tries to write to an object in The Biggest Toy Company's company data set, he meets both conditions of the CW-simple security property (because he has only read objects in that company data set), and all unsanitized objects that he can read are in the same company data set as the object he can read. Thus, both conditions of the CW-*-property are met, so Skyler can write the object.

The conditions of the CW-*-property are very restrictive; effectively, a subject can write to an object only if it has access to the company data set containing that object, *and no other company data set except the company data set in the sanitized class*. But without this restriction, conflicts of interest are possible.

9.4.5 Summary. The four models discussed in this section have played critical roles in the development of our understanding of computer security. Although it is not the first model of confidentiality, the Bell-LaPadula model describes a widely used security scheme. The Biba model captured notions of "trust" and "trustworthiness" in an intuitive way, and recent advances in the analysis of programs for vulnerabilities have applied that model to great effect. The Clark-Wilson model moved the notion of commercial integrity models away from multilevel models to models that examine process integrity as well as data integrity. The Chinese Wall model explored conflict of interest, an area that often arises when one is performing confidential services for multiple companies or has access to confidential information from a number of companies. These models are considered classic because their structure and ideas underlie the rules and structures of many other models.

9.5 OTHER MODELS. Some models examine specific environments. The Clinical Information Systems Security model[28] considers the protection of health records, emphasizing accountability as well as confidentiality and integrity. *Traducement*[29]

describes the process of real estate recordation, which requires a strict definition of integrity and accountability with little to no confidentiality.

Other models generalize the classic models. The best known are the models of *noninterference security* and *deducibility security*. Both are multilevel security models with two levels, HIGH and LOW. The noninterference model[30] defines security as the ability of a HIGH subject to interfere with what the LOW subject sees. For example, if a HIGH subject can prevent a LOW subject from acquiring a resource at a particular time, the HIGH subject can transmit information to the LOW subject. In essence, the interference is a form of writing, and must be prevented, just as the Bell-LaPadula model prevents a HIGH subject from writing to a LOW object. The deducibility model[31] examines whether a LOW subject can infer anything about a HIGH subject's actions by examining only the LOW outputs. Both these models are useful in analyzing the security of systems[32] and intrusion detection mechanisms,[33] and led to work that showed connecting two secure compute systems may produce a nonsecure system.[34] Further work is focusing on establishing conditions under which connecting two secure systems produces a secure system.[35]

9.6 CONCLUSION. The efficacy of mathematical modeling depends on the application of those models. Typically, the models capture system-specific details and describe constraints ensuring the security of the system or the information on the system. If the model does not correctly capture the details of the *entire* system, the results may not be comprehensive, and the analysis may miss ways in which security could be compromised.

This is an important point. For example, the Bell-LaPadula model captures a notion of what the *system* must do to prevent a subject cleared for TOP SECRET leaking information to a subject cleared for CONFIDENTIAL. But if the system enforces that model, the TOP SECRET subject could still meet the CONFIDENTIAL subject and hand her a printed version of the TOP SECRET information. That is outside the system and so was not captured by the model. But if the model also embraces procedures, then a procedure is necessary to prevent this "writing down." In that case, the flaw would be in the implementation of the procedure that failed to prevent the transfer of information—in other words, an incorrect instantiation of the model, exactly what Dorothy Denning's comment in the introduction to this section referred to.

The models described in this section span the foundational (access-control matrix model) to the applied (Bell-LaPadula, Biba, Clark-Wilson, and Chinese Wall). All play a role in deepening our understanding of what security is and how to enforce it.

The area of mathematical modeling is a rich and important area. It provides a basis for demonstrating that the design of systems is secure, for specific definitions of *secure*. Without these models, our understanding of how to secure systems would be diminished.

9.7 FURTHER READING

Anderson, R. "A Security Policy Model for Clinical Information Systems," Proceedings of the 1996 IEEE Symposium on Security and Privacy (May 1996): 34–48.

Bell, D., and LaPadula, L. "Secure Computer Systems: Unified Exposition and Multics Interpretation," Technical Report MTR-2997 rev. 1, MITRE Corporation, Bedford, MA (March 1975).

Biba, K. "Integrity Considerations for Secure Computer Systems," Technical Report MTR-3153, MITRE Corporation, Bedford, MA (April 1977).

Bishop, M. *Computer Security: Art and Science*. Boston: Addison-Wesley Professional, 2002.

Brewer, D., and M. Nash. "The Chinese Wall Security Policy," *Proceedings of the 1989 IEEE Symposium on Security and Privacy* (May 1989): 206–212.

Clark, D., and D. Wilson. "A Comparison of Commercial and Military Security Policies," Proceedings of the 1987 IEEE Symposium on Security and Privacy (April 1987): 184–194.

Cortier, V., and S. Kremer. *Formal Models and Techniques for Analyzing Security Protocols*—Volume 5 *Cryptology and Information Security* Series. IOS Press, 2011.

Cremers, C., and S. Mauw. *Operational Semantics and Verification of Security Protocols.* Springer, 2012.

Demillo, D., D. Dobkin, A. Jones, and R. Lipton, eds. *Foundations of Secure Computing*. New York: Academic Press, 1978.

Denning, D. "The Limits of Formal Security Models," National Information Systems Security Conference, October 18, 1999; available at www.cs.georgetown.edu/~denning/infosec/award.html

Denning, P. " Third Generation Computer Systems," *Computing Surveys* 3, No. 4 (December 1976): 175–216.

Engeler, E. *Introduction to the Theory of Computation*. New York: Academic Press, 1973.

Ferraiolo, D., J. Cugini, and D. Kuhn. "Role-Based Access Control (RBAC): Features and Motivations," Proceedings of the Eleventh Annual Computer Security Applications Conference (December 1995): 241–248.

Gougen, J., and J. Meseguer. "Security Policies and Security Models," Proceedings of the 1982 IEEE Symposium on Privacy and Security (April 1982): 11–20.

Graubert, R. "On the Need for a Third Form of Access Control," Proceedings of the Twelfth National Computer Security Conference (October 1989): 296–304.

Haigh, J., R. Kemmerer, J. McHugh, and W. Young. " An Experience Using Two Covert Channel Analysis Techniques on a Real System Design," *IEEE Transactions in Software Engineering* 13, No. 2 (February 1987): 141–150.

Harrison, M., and W. Ruzzo, " Monotonic Protection Systems." In D. Demillo et al., eds. *Foundations of Secure Computing*, pp. 337–363. New York: Academic Press, 1978.

Harrison, M., W. Ruzzo, and J. Ullman. " Protection in Operating Systems," *Communications of the ACM* 19, No. 8 (August 1976): 461–471.

Ko, C., and T. Redmond. "Noninterference and Intrusion Detection," Proceedings of the 2002 IEEE Symposium on Security and Privacy (May 2002): 177–187.

Lampson, B. "Protection." *Proceedings of the Fifth Princeton Symposium of Information Science and Systems* (March 1971): 437–443.

Lipner, S. "Non-Discretionary Controls for Commercial Applications," *Proceedings of the 1982 IEEE Symposium on Privacy and Security* (April 1982): 2–10.

Mantel, H. " On the Composition of Secure Systems," *Proceedings of the 2002 IEEE Symposium on Security and Privacy* (May 2002): 88–101.

McCullough, D. "Non-Interference and the Composability of Security Properties," *Proceedings of the 1987 IEEE Symposium on Privacy and Security* (April 1988): 177–186.

Sandhu, R. " The Typed Access Matrix Model," *Proceedings of the 1992 IEEE Symposium on Security and Privacy* (April 1992): 122–136.

Seacord, R. *Secure Coding in C and C++*. Boston: Addison-Wesley, 2005.
Walcott, T., and M. Bishop. " Traducement: A Model for Record Security," *ACM Transactions on Information Systems Security* 7, No. 4 (November 2004): 576–590.

9.8 NOTES

1. Recordation of real estate refers to recording deeds, mortgages, and other information about property with the county recorder. See http://ag.ca.gov/erds1/index.php
2. D. Denning, "The Limits of Formal Security Models," National Information Systems Security Conference, October 18, 1999; available at www.cs.georgetown.edu/~denning/infosec/award.html
3. B. Lampson, "Protection," *Proceedings of the Fifth Princeton Symposium of Information Science and Systems* (March 1971): 437–443; P. Denning, "Third Generation Computer Systems," *Computing Surveys* 3, No. 4 (December 1976): 175–216.
4. Harrison, 1976.
5. M. Harrison, W. Ruzzo, and J. Ullman, "Protection in Operating Systems," *Communications of the ACM* 19, No. 8 (August 1976): 461–471.
6. The halting problem is the question "Is there an algorithm to determine whether any arbitrary program halts?" The answer, "No," was proved by Alan Turing in 1936. See www.nist.gov/dads/HTML/haltingProblem.html. See also E. Engeler, *Introduction to the Theory of Computation* (New York: Academic Press, 1973).
7. E. Engeler, *Introduction to the Theory of Computation* (New York: Academic Press, 1973).
8. The interested reader is referred to Harrison et al., "Protection in Operating Systems," or to M. Bishop, *Computer Security: Art and Science* (Boston: Addison-Wesley Professional, 2002), p. 47 ff., for the proof.
9. Harrison et al., "Protection in Operating Systems."
10. M. Harrison and W. Ruzzo, "Monotonic Protection Systems," in D. Demillo et al., eds., *Foundations of Secure Computing,* pp. 337–363 (New York: Academic Press, 1978).
11. Harrison and Ruzzo, "Monotonic Protection Systems."
12. Harrison and Ruzzo, "Monotonic Protection Systems."
13. Harrison and Ruzzo, "Monotonic Protection Systems."
14. R. Sandhu, "The Typed Access Matrix Model," *Proceedings of the 1992 IEEE Symposium on Security and Privacy* (April 1992): 122–136.
15. Sandhu, "The Typed Access Matrix Model."
16. Also sometimes called *organization-controlled access control*, or ORGCON.
17. R. Graubert, "On the Need for a Third Form of Access Control," *Proceedings of the Twelfth National Computer Security Conference* (October 1989): 296–304.
18. D. Ferraiolo, J. Cugini, and D. Kuhn, "Role-Based Access Control (RBAC): Features and Motivations," *Proceedings of the Eleventh Annual Computer Security Applications Conference* (December 1995): 241–248.
19. As Benjamin Franklin once said, "Three can keep a secret if two of them are dead."
20. Saltzer and Schroeder 1975.
21. D. Bell and L. LaPadula, "Secure Computer Systems: Unified Exposition and Multics Interpretation," Technical Report MTR-2997 rev. 1, MITRE Corporation, Bedford, MA (March 1975).

22. A macro virus can infect a document. See, for example, Bishop, M., *Computer Security: Art and Science*. (Boston, MA: Addison-Wesley Professional, 2002), section 22.3.8, and Chapter 16 in this *Handbook*.
23. Biba proposed three models: the Low-Water-Mark Policy model, the Ring Policy model, and the Strict Integrity Policy model. See Bishop, *Computer Security*, Section 6.2.
24. The explanation is too complex to go into here. The interested reader is referred to R. Seacord, *Secure Coding in C and C++* (Boston: Addison-Wesley, 2005), Chapter 6, for a discussion of this problem.
25. S. Lipner, "Non-Discretionary Controls for Commercial Applications," *Proceedings of the 1982 IEEE Symposium on Privacy and Security* (April 1982): 2–10.
26. D. Clark and D. Wilson, "A Comparison of Commercial and Military Security Policies," *Proceedings of the 1987 IEEE Symposium on Security and Privacy* (April 1987): 184–194.
27. D. Brewer and M. Nash, "The Chinese Wall Security Policy," *Proceedings of the 1989 IEEE Symposium on Security and Privacy* (May 1989): 206–212.
28. R. Anderson, "A Security Policy Model for Clinical Information Systems," *Proceedings of the 1996 IEEE Symposium on Security and Privacy* (May 1996): 34–48.
29. T. Walcott and M. Bishop, "Traducement: A Model for Record Security," *ACM Transactions on Information Systems Security* 7, No. 4 (November 2004): 576–590.
30. J. Gougen and J. Meseguer, "Security Policies and Security Models," *Proceedings of the 1982 IEEE Symposium on Privacy and Security* (April 1982): 11–20.
31. Gougen and Meseguer, "Security Policies and Security Models."
32. J. Haigh, R. Kemmerer, J. McHugh, and W. Young, "An Experience Using Two Covert Channel Analysis Techniques on a Real System Design," *IEEE Transactions in Software Engineering* 13, No. 2 (February 1987): 141–150.
33. C. Ko and T. Redmond, "Noninterference and Intrusion Detection." *Proceedings of the 2002 IEEE Symposium on Security and Privacy* (May 2002): 177–187.
34. D. McCullough, "Non-Interference and the Composability of Security Properties," *Proceedings of the 1987 IEEE Symposium on Privacy and Security* (April 1988): 177–186.
35. H. Mantel, "On the Composition of Secure Systems," *Proceedings of the 2002 IEEE Symposium on Security and Privacy* (May 2002): 88–101.

CHAPTER 10

UNDERSTANDING STUDIES AND SURVEYS OF COMPUTER CRIME

M. E. Kabay

10.1 INTRODUCTION **10·1**
- 10.1.1 Value of Statistical Knowledge Base 10·1
- 10.1.2 Limitations on Our Knowledge of Computer Crime 10·2
- 10.1.3 Limitations on the Applicability of Computer Crime Statistics 10·2

10.2 BASIC RESEARCH METHODOLOGY **10·3**
- 10.2.1 Some Fundamentals of Statistical Design and Analysis 10·3
- 10.2.2 Research Methods Applicable to Computer Crime 10·9

10.3 SUMMARY **10·11**

10.4 FURTHER READING **10·12**

10.5 NOTES **10·12**

10.1 INTRODUCTION. This chapter provides guidance for critical reading of research results about computer crime. It will also alert designers of research instruments who may lack formal training in survey design and analysis to the need for professional support in developing questionnaires and analyzing results.

10.1.1 Value of Statistical Knowledge Base. Security specialists are often asked about computer crime; for example, customers want to know who is attacking which systems, how often, using what methods. These questions are perceived as important because they bear on the strategies of risk management; in theory, in order to estimate the appropriate level of investment in security, it would be helpful to have a sound grasp of the probability of different levels of damage. Ideally, one would want to evaluate an organization's level of risk by evaluating the experiences of other organizations with similar system and business characteristics. Such comparisons would be useful in competitive analysis and in litigation over standards of due care and diligence in protecting corporate assets.

10.1.2 Limitations on Our Knowledge of Computer Crime. Unfortunately, in the current state of information security, no one can give reliable answers to such questions. There are two fundamental difficulties preventing us from developing accurate statistics of this kind. These difficulties are known as the problems of ascertainment.

10.1.2.1 Detection. The first problem is that an unknown number of crimes of all kinds are undetected. For example, even outside the computer crime field, we do not know how many financial frauds are being perpetrated. We do not know because some of them are not detected. How do we know they are not detected? Because some frauds are discovered long after they have occurred. Similarly, computer crimes may not be detected by their victims but may be reported by the perpetrators.

In a landmark series of tests at the Department of Defense (DoD), the Defense Information Systems Agency found that very few of the penetrations it engineered against unclassified systems within the DoD seem to have been detected by system managers. These studies were carried out from 1994 through 1996 and attacked 38,000 systems. About two-thirds of the attacks succeeded; however, only 4 percent of these attacks were detected.[1]

A commonly held view within the information security community is that only one-tenth or so of all the crimes committed against and using computer systems are detected.

10.1.2.2 Reporting. The second problem of ascertainment is that even if attacks are detected, few seem to be reported in a way that allows systematic data collection. This commonly held belief is based in part on the unquantified experience of information security professionals who have conducted interviews of their clients; it turns out that only about 10 percent of the attacks against computer systems revealed in such interviews were ever reported to any kind of authority or to the public. The DoD studies mentioned earlier were consistent with this belief; of the few penetrations detected, only a fraction of 1 percent were reported to appropriate authorities.

Given these problems of ascertainment, computer crime statistics generally should be treated with skepticism.

10.1.3 Limitations on the Applicability of Computer Crime Statistics. Generalizations in this field are difficult to justify. Even if we knew more about types of criminals and the methods they use, it still would be difficult to have the kind of actuarial statistic that is commonplace in the insurance field. For example, the establishment of uniform building codes in the 1930s in the United States led to the growth in fire insurance as a viable business. With official records of fires in buildings that could be described using a standard typology, statistical information began to provide an actuarial basis for using probabilities of fires and associated costs to calculate reasonable insurance rates.

In contrast, even if we had access to accurate reports, it would be difficult to make meaningful generalizations about vulnerabilities and incidence of successful attacks for the information technology field. We use a bewildering variety and versions of processors, operating systems, firewalls, encryption, application software, backup methods and media, communications channels, identification, authentication, authorization, compartmentalization, and operations.

How would we generalize from data about the risks at (say) a mainframe-based network running Multiple Virtual Systems (MVS) in a military installation to the kinds of risks faced by a UNIX-based intranet in an industrial corporation, or to a Windows New Technology (NT)–based Web server in a university setting? There are so many differences among systems that if we were to establish a multidimensional analytical table where every variable was an axis, many cells would likely contain no or only a few examples. Such sparse matrices are notoriously difficult to use in building statistical models for predictive purposes.

10.2 BASIC RESEARCH METHODOLOGY. This is not a chapter about social sciences research. However, many discussions of computer crime seem to take published reports as gospel, even though these studies may have no validity whatsoever. In this short section, we look at some fundamentals of research design so that readers will be able to judge how much faith to put in computer crime research results.

10.2.1 Some Fundamentals of Statistical Design and Analysis. The way in which a scientist or reporter represents data can make an enormous difference in the readers' impressions.

10.2.1.1 Descriptive Statistics. Suppose three companies reported these losses from penetration of their computer systems: $1 million, $2 million, and $6 million. We can describe these results in many ways. For example, we can simply list the raw data; however, such lists could become unacceptably long as the number of reports increased, and it is hard to make sense of the raw data.

We could define classes such as "2 million or less" and "more than 2 million" and count how many occurrences there were in each class:

Class	Freq
≤ $2M	2
> $2M	1

Alternatively, we might define the classes with finer granularity as < $1M, ≥ $1M but < $2M, and so on; such a table might look like this:

Class	Freq
< $1M	0
≥ $1M & < $2M	1
≥ $2M & < $3M	1
≥ $3M & < $4M	0
≥ $4M & < $5M	0
≥ $5M & < $6M	0
≥ $6	1

Notice how the definition of the classes affects perception of the results: The first table gives the impression that the results are clustered around $2 million and gives no information about the upper or lower bounds.

10.2.1.1.1 Location. One of the most obvious ways we describe data is to say where they lie in a particular dimension. The *central tendency* of our three original data ($1 million, $2 million, and $6 million) can be represented in various ways; for example, two popular measures are:

- Arithmetic mean or average = $(1+2+6)M/3 = $3M
- Median (the middle of the sorted list of losses) = $2M

Note that if we tried to compute the mean and the median from the first table (with its approximate classes), we would get the wrong value. Such statistics should be computed from the original data, not from summary tables.

10.2.1.1.2 Dispersion. Another aspect of our data that we frequently need is *dispersion*—that is, variability. The simplest measure of dispersion is the range: the difference between the smallest and the largest value we found; in our example, we could say that the range was from $1 million to $6 million or that it was $5 million. Sometimes the range is expressed as a percentage of the mean; then we would say that the range was $5/3 = 1.6 \ldots$ or ~167 percent.

The *variance* (σ^2) of these particular data is the average of the squared deviations from the arithmetic mean; the variance of the three numbers would be $\sigma^2 = (1-3)^2 + (2-3)^2 + (6-3)^2/3 = (4+1+9)/3 \approx 4.67$.

The square root of the variance (σ) is called the *standard deviation* and is often used to describe dispersion. In our example, $\sigma = \sqrt{4.67} \approx 2.16$.

Dispersion is particularly important when we compare estimates about information from different groups. The greater the variance of a measure, the more difficult it is to form reliable generalizations about an underlying phenomenon, as described in the next section.

10.2.1.2 Inference: Sample Statistics versus Population Statistics. We can accurately describe any data using descriptive statistics; the question is what we then do with those measures.

Usually we expect to extend the findings in a *sample* or subset of a *population* to make generalizations about the population. For example, we might be trying to estimate the losses from computer crime in commercial organizations with offices in the United States and with more than 30,000 employees. Or perhaps our sample would represent commercial organizations with offices in the United States and with more than 30,000 employees and whose network security staff was willing to respond to a survey questionnaire.

In such cases, we try to infer the characteristics of the population from the characteristics of the sample. Statisticians say that we try to estimate the *parametric* statistics by using the *sample* statistics.

For example, we estimate the parametric (population) variance (usually designated σ^2) by multiplying the variance of the sample by $n/(n-1)$. Thus, we would say that the estimate of the parametric variance (s^2) in our sample would be $s^2 = 4.67 * 3/2 = 7$. The estimate of the parametric standard deviation (s) would be $s = \sqrt{7} \approx 2.65$.

10.2.1.3 Hypothesis Testing. Another kind of inference that we try to make from data is *hypothesis testing*. For example, suppose we were interested in whether there was any association between the presence or absence of firewalls and the occurrence of system penetration. We can imagine collecting these data about penetrations into systems with or without firewalls:

	Penetration		
Firewalls	***No***	***Yes***	***Totals***
No	25	75	100
Yes	70	130	200
Totals	95	205	300

We would frame the hypothesis (the *null hypothesis,* sometimes represented as H_0) that there was *no* relationship between the two independent variables, penetration and firewalls, and *test* that hypothesis by performing a test of independence of these variables. In our example, a simple chi-square test of independence would give a *test statistic* of $\chi^2_{[1]} = 2.636$. If there really were no association between penetration and firewalls in the population of systems under examination, the parametric value of this statistic would be zero. In our imaginary example, we can show that such a large value (or larger) of $\chi^2_{[1]}$ would occur in only 10.4 percent of the samples taken from a population where firewalls had no effect on penetration. Put another way, if we took many samples from a population where the presence of firewalls was not associated with any change in the rate of penetration, we would see about 10.4 percent of those samples producing $\chi^2_{[1]}$ statistics as large as or larger than 2.636.

Statisticians have agreed on some conventions for deciding whether a test statistic deviates enough from the value expected under the null hypothesis to warrant inferring that the null hypothesis is wrong. Generally, we describe the likelihood that the null hypothesis is true—often shown as $p(H_0)$—in this way:

- When $p(H_0) > 0.05$, we say the results are *not statistically significant* (often designated with the symbols *ns*);
- When $0.05 \geq p(H_0) > 0.01$, the results are described as statistically significant (often designated with the symbol *);
- When $0.01 \geq p(H_0) > 0.001$, the results are described as highly statistically significant (often designated with the symbols **);
- When $p(H_0) \leq 0.001$, the results are described as extremely statistically significant (often designated with the symbols ***).

10.2.1.4 Random Sampling, Bias, and Confounded Variables. The most important element of sampling is randomness. We say that a sample is *random* or *randomized* when every member of the population we are studying has an equal probability of being selected. When a population is defined one way but the sample is drawn nonrandomly, the sample is described as *biased.* For example, if the population we are studying was designed to be, say, all companies worldwide with more than 30,000 full-time employees, but we sampled mostly from such companies in the United States, the sample would be biased toward U.S. companies and their characteristics.

Similarly, if we were supposed to be studying security in all companies in the United States with more than 30,000 full-time employees, but we sampled only from those companies that were willing to respond to a security survey, we would be at risk of having a biased sample.

In this last example, involving studying only those who respond to a survey, we say that we are potentially *confounding* variables: We are looking at people-who-respond-to-surveys and hoping they are representative of the larger population of people from all companies in the desired population. But what if the people who are *willing* to respond are those who have better security and those who do not respond have terrible security? Then *responding to the survey* is confounded with *quality of security,* and our biased sample could easily mislead us into overestimating the level of security in the desired population.

Another example of how variables can be confounded is comparisons of results from surveys carried out in different years. Unless exactly the same people are interviewed in both years, we may be confounding individual variations in responses with changes over time; unless exactly the same companies are represented, we may be confounding differences among companies with changes over time; if external events have led people to be more or less willing to respond truthfully to questions, we may be confounding willingness to respond with changes over time. If the surveys are carried out with different questions or used by different research groups, we may be confounding changes in methodology with changes over time.

10.2.1.5 Confidence Limits. Because random samples naturally vary around the parametric (population) statistics, it is not very helpful to report a *point estimate* of the parametric value. For example, if we read that the mean damage from computer crimes in a survey was $180,000 per incident, what does that imply about the population mean?

To express our confidence in the sample statistic, we calculate the likelihood of being right if we give an *interval estimate* of the population value. For example, we might find that we would have a 95 percent likelihood of being right in asserting that the mean damage was between $160,000 and $200,000. In another sample, we might be able to narrow these *95 percent confidence limits* to $175,000 and $185,000.

In general, the larger the sample size, the narrower the confidence limits will be for particular statistics.

The calculation of confidence limits for statistics depends on some necessary *assumptions*:

- Random sampling
- A *known error distribution* (usually the *Normal* distribution—sometimes called a *Gaussian* distribution)
- *Equal variance* at all values of the measurements

If any of these assumptions is wrong, the calculated confidence limits for our estimates will be wrong; that is, they will be misleading. There are tests of these assumptions that analysts should carry out before reporting results; if the data do not follow Normal error distributions, sometimes one can apply *normalizing transformations*.

In particular, percentages do not follow a Normal distribution. Here is a reference table of confidence limits for various percentages in a few representative sample sizes.

95 Percent Confidence Limits for Percentages			
	Sample size		
Percentage	100	500	1000
0	0–3.0%	0–0.6%	0–0.3%
10	4.9–17.6%	7.5–13.0%	8.2–12.0%
20	12.7–29.1%	16.6–23.8%	17.6–22.6%
50	40.0–60.1%	45.5–54.5%	46.9–53.1%
80	70.9–87.3%	76.2–83.4%	77.4–82.4%
90	82.4–95.1%	87.0–92.5%	88.0–91.8%
100	97.0–100%	99.4–100%	99.7–100%

10.2.1.6 Contingency Tables. One of the most frequent errors in reporting results of studies is to provide only part of the story. For example, one can read statements such as "Over 70 percent of the systems without firewalls were penetrated last year." Such a statement may be true, but it cannot be interpreted correctly as meaning that systems with firewalls were necessarily more or less vulnerable to penetration than systems without firewalls. The statement is incomplete; to make sense of it, we need the other part of the implied *contingency table*—the percentage of systems *with* firewalls that were penetrated last year—before making any assertions about the relationship between firewalls and penetrations. Compare, for example, these two hypothetical tables:

	Without Firewalls	With Firewalls in Default Configuration
Penetrated	70%	70%
Not Penetrated	30%	30%

	Without Firewalls	With Firewalls Properly Configured
Penetrated	70%	10%
Not Penetrated	30%	90%

In both cases, someone could say that "70 percent of the systems without firewalls were penetrated," but the implications would be radically different in the two data sets. Without knowing the right-hand column, the original assertion would be meaningless.

10.2.1.7 Association versus Causality. Continuing our example with rates of penetration, another error that untrained people often make when studying statistical information is to mistake *association* for *causality*. Imagine that a study showed that a lower percentage of systems with fire extinguishers was penetrated than systems without fire extinguishers and that this difference was statistically highly significant. Would such a result necessarily mean that fire extinguishers *caused* the reduction in penetration? No. We know that it is far more reasonable to suppose that the fire extinguishers were installed in organizations whose security awareness and security policies were more highly developed than in the organizations where no fire extinguishers were installed. In this imaginary example, the fire extinguishers might actually have *no causal effect* whatever on resistance to penetration. This result would illustrate the effect of

confounding variables: *presence of a fire extinguisher* with *state of security awareness and policies*.

10.2.1.8 Control Groups. Finally, to finish our penetration example, one way to distinguish between association and causality is to *control* for variables. For example, one could measure the state of security awareness and policy as well as the presence or absence of fire extinguishers and make comparisons only among groups with the same level of awareness and policy. There are also statistical techniques for mathematically controlling for differences in such *independent variables*.

10.2.1.9 A Priori versus a Posteriori Testing. Amateurs or beginners sometimes forget the principle of *random sampling* that underlies all statistical inference (see Section 10.2.1.4). None of the hypothesis tests or confidence limit calculations work if a sample is not random. For example, if someone is wandering through a supermarket and notices that Granny Smith apples seem to be bigger than Macintosh apples, selecting a sample—even a random sample—of the apples that specifically gave rise to the hypothesis will not allow reliable computations of probability that the applies have the same average weight. The problem is that those *particular* apples would not have been sampled at all had the observer not been moved to formulate the hypothesis. So even if a particular statistical comparison produces a sample statistic that appears to have a probability of, say, 0.001, it is not possible to know how much the sampling deviated from randomness.

Applying statistical tests to data *after* one notices an interesting descriptive value, comparison, or trend is known as *a posteriori testing.* Formulating a hypothesis, obtaining a random sample, and computing the statistics and probabilities in accordance with the assumptions of those statistics and probabilities is known as *a priori testing*.

A well-used example of the perils of a posteriori testing is the unfortunate habit of searching through sequences of results such as long strings of guesses collected in student tests of paranormal abilities and calculating statistical values on carefully selected subsets of the strings. These a posteriori tests are then presented as if they were a priori and cause great confusion and arguments, such as: "Look, even though the overall proportion of correct guesses was (say) 50.003 percent in this run of <some very large number> guesses, there was a run of <much smaller number> guesses that were correct <any value greater than 50 percent> of the time! The probability of such a result by chance is <very small number>. That proves that there was a real effect of <whatever the treatment was>." Unfortunately, a long series of numbers can produce any desired nonrandom-looking string; there are even tests known as *runs tests* that can help a researcher evaluate the nonrandomness of such occurrences.

In practical terms, statisticians have established a convention for limiting the damaging effects of a posteriori testing: Use the 0.001 level of probability as the equivalent of the minimum probability of the null hypothesis. This custom makes it far less likely that an a posteriori comparison will trick the user into accepting what is in fact a random variation that caught someone's eye.

The best solution to the bias implicit in a posteriori testing is to use a completely *new* sample for the comparison. In the apple example, one could ask the store manager for new, unobserved, and randomly selected batches of both types of applies. The comparison statistics would then be credible and could be expected to follow the parametric distribution underlying calculations of probability of the null hypothesis. The populations from which these apples were selected still would have to be carefully

determined. Would the populations be apples at this particular store? For this particular chain? For this particular region of the country or of the world?

10.2.2 Research Methods Applicable to Computer Crime

10.2.2.1 Interviews. Interviewing individuals can be illuminating. In general, interviews provide a wealth of data that are unavailable through any other method. For example, one can learn details of computer crime cases or motivations and techniques used by computer criminals. Interviews can be structured (using precise lists of questions) or unstructured (allowing the interviewer to respond to new information by asking additional questions at will).

Interviewers can take notes or record the interviews for later word-for-word transcription. In unstructured interviewers, skilled interviewers can probe responses to elucidate nuances of meaning that might be lost using cruder techniques such as surveys. Techniques such as thematic analysis can reveal patterns of responses that can then be examined using exploratory data analysis.[2] Thematic analysis is a technique for organizing nonquantitative information without imposing a preexisting framework on the data; exploratory data analysis uses statistical techniques to identify possibly interesting relationships that can be tested with independently acquired data. Such exploratory techniques can correctly include a posteriori testing as described in Section 10.2.1.9, but the results are used to propose further studies that can use a priori tests for the best use of resources.

10.2.2.2 Focus Groups. Focus groups are like group interviews. Generally, the facilitator uses a list of predetermined questions and encourages the participants to respond freely and to interact with each other. Often, the proceedings are filmed from behind a one-way mirror for later detailed analysis. Such analysis can include nonverbal communications, such as facial expressions and other body language as the participants speak or listen to others speak about specific topics.

10.2.2.3 Surveys. Surveys consist of asking people to answer a fixed series of questions with lists of allowable answers. They can be carried out face to face or by distributing and retrieving questionnaires by telephone, mail, fax, and email. Some questionnaires have been posted on the Web.

The critical issue when considering the reliability of surveys is *self-selection bias*—the obvious problem that survey results include only the responses of people who agreed to participate. Before basing critical decisions on survey data, it is useful to find out what the response rate was; although there are no absolutes, in general, we tend to trust survey results more when the response rate is high. Unfortunately, response rates for telephone surveys are often less than 10 percent; response rates for mail and email surveys can be less than 1 percent. It is very difficult to make any case for random sampling under such circumstances, and all results from such low-response-rate surveys should be viewed as indicating the range of problems or experiences of the respondents rather than as indicators of population statistics.

Regarding Web-based surveys, there are two types from a statistical point of view: those that use strong identification and authentication and those that do not. Those that do not are vulnerable to fraud, such as repeated voting by the same individuals. Those that provide individual universal resource locators (URLs) to limit voting to one per

person nonetheless suffer from the same problems of self-selection bias as any other survey.

10.2.2.4 Instrument Validation. Interviews and other social sciences research methodologies can suffer from a systematic tendency for respondents to shape their answers to please the interviewer or to express opinions that may be closer to the norm in whatever group they see themselves. Thus, if it is well known that every organization ought to have a business continuity plan, some respondents may misrepresent the state of their business continuity planning to look better than they really are.

In addition, survey instruments may distort responses by phrasing questions in a biased way; for example, the question "Does your business have a completed business continuity plan?" may have a more accurate response rate than the question "Does your business comply with industry standards for having a completed business continuity plan?" The latter question is not neutral and is likely to increase the proportion of "yes" answers.

The sequence of answers may bias responses; exposure to the first possible answers can inadvertently establish a baseline for the respondent. For example, a question about the magnitude of virus infections might ask:

In the last 12 months, has your organization experienced total losses from virus infections of

(a) $1 million or greater;
(b) less than $1 million but greater than or equal to $100,000;
(c) less than $100,000;
(d) none at all?

To test for bias, the designer can create versions of the instrument in which the same information is obtained using the opposite sequence of answers:

In the last 12 months, has your organization experienced total losses from virus infections of

(a) none at all;
(b) less than $100,000;
(c) less than $1 million but greater than or equal to $100,000;
(d) $1 million or greater?

The sequence of questions can bias responses; having provided a particular response to a question, the respondent will tend to make answers to subsequent questions about the same topic conform to the first answer in the series. To test for this kind of bias, the designer can create versions of the instrument with questions in different sequences.

Another instrument validation technique inserts questions with no valid answers or with meaningless jargon to see if respondents are thinking critically about each question or merely providing any answer that pops into their heads. For example, one might insert the nonsensical question, "Does your company use steady-state quantum interference methodologies for intrusion detection?" into a questionnaire about security and invalidate the results of respondents who answer yes to this and other diagnostic questions.

Finally, independent verification of answers provides strong evidence of whether respondents are answering truthfully. However, such intrusive investigations are rare.

10.2.2.5 Meta-analysis. Sometimes it is useful to evaluate a hypothesis based on several studies. We can combine probabilities *P* of the *same* null hypothesis from *k* trials using the formula

$$X^2 = -2\Sigma \ln P$$

where X^2 is distributed as $\chi^2_{[2k]}$ if the null hypothesis is true in all the trials.

For example, suppose a forensic specialist is evaluating the possibility that log files have been tampered with in a specific period.

- She runs a test of the frequency distribution of individual digits in the data for the suspect period to evaluate the likelihood that the distribution is consistent with the null hypothesis of randomness. The P for the two-tailed chi-square test is 0.072ns.
- She looks at the average number of disk I/Os per second in the records for the suspect period and compares the data with the same statistic in the control period using ANOVA; the P for the two-tailed test is 0.096ns.

In this case, $X^2 = -2\ \Sigma \ln P\ X^2 = -2^*(\ln 0.072 + \ln 0.096) = -2^*(-2.63109 -2.34341) = 9.948992$ with 4 degrees of freedom. The P for the null hypothesis is 0.0413^*. In other words, the chances of observing the results of both tests by chance alone if the suspect data were consistent with the raw data from comparison data is statistically significant. There is reason to reject the null hypothesis: Someone may very well have tampered with the log files for the suspect period.

This technique is subject to constraints. Most meta-analyses require the probabilities of the null hypothesis to be computed for a single tail in the same direction; it doesn't make sense to combine probabilities from conflicting hypotheses using two-tailed probabilities. However, as in the example above, if the null hypotheses are consistent, even two-tailed probabilities may be usefully combined.

Another problem is more systemic. There is considerable reason to be concerned that investigators and publishers sometimes suppress experimental results that do not conform to their expectations or desires. Such suppression biases the published results to *appear* to support the desired result.

Identifying such suppression is difficult. One approach is to examine the distributions of the published data and look for indications of data exclusion. For example, if the frequency distribution for raw data in a published report shows an abrupt disappearance of data in one direction (e.g., if the frequency distribution looks like a normal curve except that the left side suddenly drops to zero at a certain point) then the publication may be using a truncated data set. Meta-analysis based on data of dubious validity will itself be dubious. Garbage in, garbage out.

10.3 SUMMARY. In summary, all studies about computer crime should be studied carefully before we place reliance on their results. Some basic take-home questions about such research:

- What is the population we are sampling?
- Keeping in mind the self-selection bias, how representative of the wider population are the respondents who agreed to participate in the study or survey?
- How large is the sample?
- Are the authors testing for the assumptions of randomness, normality, and equality of variance before reporting statistical measures?

- What are the confidence intervals for the statistics being reported?
- Are comparisons confounding variables?
- Are correlations being misinterpreted as causal relations?
- Were the test instruments validated?

10.4 FURTHER READING

Textbooks

If you are interested in learning more about survey design and statistical methods, you can study any elementary textbook on the social sciences statistics. Here are some sample titles:

Babbie, E. R., F. S. Halley, and J. Zaino. *Adventures in Social Research: Data Analysis Using SPSS 11.0/11.5 for Windows*, 5th ed. Pine Forge Press, 2003.

Bachman, R., and R. K. Schutt. *The Practice of Research in Criminology and Criminal Justice*, 3rd ed. Sage Publications, 2007.

Carlberg, C. *Statistical Analysis: Microsoft Excel 2010*. Que, 2011.

Chambliss, D. F., and R. K. Schutt. *Making Sense of the Social World: Methods of Investigation*, 2nd ed. Pine Forge Press, 2006.

Cox, D. R., and C. A. Donnelly. *Principles of Applied Statistics*. Cambridge University Press (ISBN 978-1107644458), 2011. 212 pp.

Kabay, M. E. *Statistics in Business, Finance, Management, and Information Technology: A Layered Introduction with Excel*. Free textbook (PDF), 2013. 205 pp. www.mekabay.com/courses/academic/norwich/qm213/statistics_text.pdf

Sirkin, R. M. *Statistics for the Social Sciences*, 3rd ed. Sage Publications, 2005.

Warner, R. M. *Applied Statistics: From Bivariate through Multivariate Techniques*, 2nd ed. SAGE Publications (ISBN 978-1412991346), 2012. 1208 pp.

Websites

Education Insider (2011) "Explore Statistics in the Blogosphere: Top 10 Statistics Blogs." http://education-portal.com/articles/Explore_Statistics_in_the_Blogosphere_Top_10_Statistics_Blogs.html

StatPac, "Survey & Questionnaire Design," www.statpac.com/surveys

10.5 NOTES

1. GAO (1996) "Computer Attacks at Department of Defense Pose Increasing Risks." General Accounting Office Report to Congressional Requesters GAO/AIMD-98-84 (May 1996), p. 19
2. M. E. Kabay, "CATA: Computer-Aided Thematic Analysis," 2006. Available at www.mekabay.com with narrated lectures at www.mekabay.com

CHAPTER 11

FUNDAMENTALS OF INTELLECTUAL PROPERTY LAW

William A. Zucker and Scott J. Nathan

11.1 INTRODUCTION 11·2

11.2 THE MOST FUNDAMENTAL BUSINESS TOOL FOR PROTECTION OF TECHNOLOGY IS THE CONTRACT 11·3

11.2.1 Prevention Begins at Home—Employee and Fiduciary Duties 11·4

11.2.2 Employment Contract, Manual, and Handbook 11·4

11.2.3 Technology Rights and Access in Contracts with Vendors and Users 11·4

11.3 PROPRIETARY RIGHTS AND TRADE SECRETS 11·5

11.3.1 Remedies for Trade Secret Misappropriation 11·6

11.3.2 Vigilance Is a Best Practice 11·8

11.4 COPYRIGHT LAW AND SOFTWARE 11·8

11.4.1 Works for Hire and Copyright Ownership 11·9

11.4.2 Copyright Rights Adhere from the Creation of the Work 11·9

11.4.3 First Sale Limitation 11·9

11.4.4 Fair Use Exception 11·10

11.4.5 Formulas Cannot Be Copyrighted 11·10

11.4.6 Copyright Does Not Protect the "Look and Feel" for Software Products 11·10

11.4.7 Reverse Engineering as a Copyright Exception 11·11

11.4.8 Interfaces 11·11

11.4.9 Transformative Uses 11·12

11.4.10 Derivative Works 11·12

11.4.11 Semiconductor Chip Protection Act of 1984 11·13

11.4.12 Direct, Contributory, or Vicarious Infringement 11·13

11.4.13 Civil and Criminal Remedies 11·13

11.5 DIGITAL MILLENNIUM COPYRIGHT ACT 11·14

11.6 CIRCUMVENTING TECHNOLOGY MEASURES 11·15

11.6.1 Exceptions to the Prohibitions on Technology Circumvention 11·16

11.7 PATENT PROTECTION 11·18

11.7.1 Patent Protection Requires Disclosure 11·18

11.7.2 Patent Protection in Other Jurisdictions 11·19
11.7.3 Patent Infringement 11·19

11.8 PIRACY AND OTHER INTRUSIONS 11·20
11.8.1 Marketplace 11·20
11.8.2 Database Protection 11·20
11.8.3 Applications of Transformative and Fair Use 11·21
11.8.4 Internet Hosting and File Distribution 11·21
11.8.5 Web Crawlers and Fair Use 11·23
11.8.6 HyperLinking 11·23
11.8.7 File Sharing 11·23

11.9 OTHER TOOLS TO PREVENT UNAUTHORIZED INTRUSIONS 11·24
11.9.1 Trespass 11·24
11.9.2 Terms of Use 11·25
11.9.3 Computer Fraud and Abuse Act 11·26
11.9.4 Electronic Communications and Privacy 11·29
11.9.5 Stored Communications Act 11·31

11.10 OPEN SOURCE 11·33
11.10.1 Open Source Licenses 11·33
11.10.2 GPL 11·33
11.10.3 Other Open Source Licenses 11·34
11.10.4 Business Policies with Respect to Open Source Licenses 11·34

11.11 APPLICATION INTERNATIONALLY 11·34
11.11.1 Agreement on Trade-Related Aspects of Intellectual Property Rights 11·35
11.11.2 TRIPS and Trade Secrets 11·36
11.11.3 TRIPS and Copyright 11·37
11.11.4 TRIPS and Patents 11·37
11.11.5 TRIPS and Anticompetitive Restrictions 11·38
11.11.6 Remedies and Enforcement Mechanisms 11·38

11.12 RECENT DEVELOPMENTS IN INTELLECTUAL PROPERTY LAW 11·39
11.12.1 AIA 11·39
11.12.2 The PROTECT IP Act (PIPA) 11·39
11.12.3 The Stop Online Piracy Act (SOPA) 11·41
11.12.4 Patent Trolls 11·43

11.13 CONCLUDING REMARKS 11·44

11.14 FURTHER READING 11·44

11.15 NOTES 11·44

11.1 INTRODUCTION. This chapter is not for lawyers or law students. Rather, it is written for computer professionals who might find it useful to understand how their concerns at work fit into a legal framework, and how that framework shapes strategies that they might employ in their work. It is not intended to be definitive but to help readers spot issues when they arise and to impart an understanding that is the first part of a fully integrated computer security program.

Cyberlaw is a compendium of traditional law that has been updated and applied to new technologies. When gaps have developed or traditional law is inadequate, particular statutes have been enacted. It is a little like the old story of the three blind men and the elephant: One of the blind men touching the elephant's leg believes he is touching a tree; the other touching its ear believes it is a wing, and the third, touching the tail, thinks it is a snake. Issues of cyberspace, electronic data, networks, global transmissions, and positioning have neither simple unitary solutions nor a simple body of law to consult.

In thinking about the application of law to computer security, it is helpful to think about the problems as issues in which the computer is

- The target of the activity
- The tool used for the activity
- Incidental to the activity itself

For example, "hacking" into a computer can be analogized to the tort[1] of trespass (i.e., entering the property of another without permission), and "cracking" can be viewed as conversion of someone else's property. Similarly, using the computer to make illegal copies is a violation of copyright law in its most basis sense. Although trademark law has very little to do with computers, using trade names as part of keywords for search engines, or domain names to misdirect Internet traffic to a competitive Web site can be a violation of a trademark. While touching on some of the more traditional tort remedies, this chapter focuses on the property rights being invaded by such activities and the remedies that exist in the context of a business operation.

Recognizing that the body of law which touches on these problems is as global as the Internet itself, this chapter is intended to help readers actually see the elephant in the room. In selecting what legal issues to highlight, we have tried to consider the routine needs of the computer professional. We have focused largely on the law of the United States, recognizing that these problems and subject matters often transcend national boundaries. There is a very simple reason for this. Most often, the impact of the computer security attack, denial of service, decryption, or theft of computer materials will have occurred here, or have a direct impact here, no matter where it originates. Imagine for a second a gunman—standing in Canada—who takes aim at someone in the United States, pulls the trigger, and hits his target. Since there is purposeful conduct aimed at this country, in the ordinary instance the U.S. judiciary will not only assert jurisdiction over the gunman but also apply its laws. There may be other problems, such as actually catching the gunman, but the example underlines the importance of the law of the United States for entities located here. For orientation purposes, we have also included a section at the end of this chapter that discusses some international issues.

One other introductory note: We use the phrase "security program" in this chapter with some frequency. Understanding that this phrase can mean one thing to a lawyer or risk manager and another thing to a computer security professional, we intend it as a shorthand reference to the generic and systemic effort to secure information stored on computers and not solely to the applications that may be employed as part of that effort.

11.2 THE MOST FUNDAMENTAL BUSINESS TOOL FOR PROTECTION OF TECHNOLOGY IS THE CONTRACT.

The computer security professional's job is to understand, anticipate, and then worry about risk: risks that are beyond control and risks that can be controlled. The most fundamental tool for controlling risk, whether predictable or unforeseeable, is the contract. Unlike other forms of risk control, a contract need not be static; it can be adaptable. We can limit use; we can limit distribution; we can impose conditions and confidentiality; we can specify rights as well as provide for certain remedies through contract. Contracts actually can take many forms: the traditional signed agreement; an email exchange; Web site or product terms of use; employment agreements; workplace manuals and policies; and so-called shrink-wrap or click-wrap agreements. We sell or license products. Where we can contract, we can also define and limit risk.

11.2.1 Prevention Begins at Home—Employee and Fiduciary Duties. There is an old hoary concept in the law that employees owe to their employers the fiduciary duty of utmost loyalty. The scope and extent of that fiduciary duty is a matter of common law that varies in each state. Generally, employees' fiduciary duty prohibits them from using any property that belongs to the employer in competition with the employer or for personal gain. Employees, however, are entitled to retain and use for whatever purpose their own skill and knowledge, which arguably could include contacts that they develop over the course of their employment unless those contacts are trade secrets. What comes or does not come within the ambit of fiduciary duty has spawned endless arguments and lawsuits. There is a simple remedy to this problem: the contract that covers technology issues and ownership as well as it covers pay and other benefits.

11.2.2 Employment Contract, Manual, and Handbook. Whatever policy the security professional develops should be implemented through the organization's employment contract, manual, and handbook. Many contractual provisions can be applied, such as: nondisclosure agreements; definition of proprietary policy; restrictive covenants; concessions of ownership regarding discoveries, know-how, improvements, inventions, and the like during the term of employment; email policies; terms of use regarding computer systems; and statements of authorized and unauthorized activity. The point is that employment contracts and handbooks should be the starting point for computer security.

11.2.3 Technology Rights and Access in Contracts with Vendors and Users. Security protection necessarily includes vigilance about all contracts and licenses with vendors and users. This may not be sexy, but it is blocking and tackling. Vendors can be subject to many of the same limitations and nondisclosure agreements as employees. Rights of access to intranets and data should be controlled and privileges specified. Careful consideration should be given to what rights a user will have, the rules surrounding user access, and enforcement of those rules. Is this a sale or a license? There are many virtues to controlling technology through licenses (as opposed to sales), including imposing limits on rights of use, and specifying remedies for breaches of the license, or for unauthorized activity that involves the licensed product.

"Shrink-wrap" or "click-wrap" licenses have become common parlance. They are now accepted tools for licensing and controlling software distribution so long as: (a) they are business to business and thus between parties of roughly equal bargaining position; (b) their terms for other users or consumers are not unconscionable; and (c) they do not violate public policy. Concerns over whether contractual terms are unconscionable or the contracts are ones of adhesion arise because the licenses are not products of negotiation but of fiat, which users accept when they open the shrink-wrapped package or through an online click. These concerns have been addressed through requirements that users have been provided with adequate notice of the terms, an opportunity to reject, and conduct that sufficiently manifests consent. For shrink-wrap agreements, the opening of the product, its installation, and retention have been deemed sufficient acts to show consent to the terms of the license, noting that if the consumer does not wish to consent, the product could be returned.[2] Thus, it is not necessary for the prospective user to be aware of all of the terms of a license before purchase if the remedy includes return after purchase. The license can impose restrictions on use, limit the number of machines on which the product can be installed, copying, and even available remedies.[3]

The issues of notice, actual or constructive, an opportunity to accept or reject, and manifestation of consent have led to general acceptance of online agreements such as the presentation of licensing terms followed by an active need to check, accept, or reject by clicking on the appropriate box.[4] The same analysis applies to terms of use especially for intranet or network use.[5] In *Register.com v. Verio, Inc.*,[6] downloading data from a WHOIS database, having knowledge of the terms of use, was acceptance of those terms even if there was no click-through. These examples show that terms of use, properly positioned, can be binding on the user.

An active security program begins with a review of the contracts, licenses, and terms of use in all relationships with your organization. Just because a contractual arrangement has not existed does not mean that you cannot create one through proper notice of the terms of the contract and conduct that shows assent to those terms. Such contracts are the security professional's first line of defense. They give you the ability to limit risk with an organization's employees, contractors, vendors, and affiliates. With that in mind, this chapter addresses issues that arise largely outside of the terms of contractual protections and also suggests additional potential self-help remedies.

11.3 PROPRIETARY RIGHTS AND TRADE SECRETS. For many years, unless an idea was patentable, the primary protection for internal business data, confidential or proprietary information, and computer code was through the common law doctrine of trade secrets.[7] Generally, a trade secret might be considered any internal, nonpublished manufacturing know-how, drawings, formulas, or sales information used in a trade or business that has commercial applicability and that provides a business with some strategic advantage.[8] Such information, so long as it was (a) not published or disseminated to others who were not obligated to maintain its confidentiality,[9] and (b) maintained in confidence with the protecting organization, could be protected as a trade secret.

The law of trade secret thus recognized a business's ownership or proprietary interest in such information, data, or processes. There are, however, important practical limitations on the application of trade secret protection. First and foremost, for any product sold in the market, the law does not protect against a competitor seeing the product and then using it to figure out how to manufacture like or similar items. Competitors are therefore free to reverse engineer a product so long as the reverse engineering is done wholly independently.

The second caveat is that an organization has to prove not only that the information qualifies for trade secret protection, but also that it protected the secrecy of the information as required by the law of the applicable jurisdiction. This means that ownership will be a matter not of record but of case-by-case proof, making enforcement of trade secret protection time consuming and expensive. Generally, the required proof consists of a showing that there was an active security program in place that was sufficient to protect the information as confidential. Various programs may be deemed adequate, depending on the circumstances, but usually such programs have five principles in common:

1. An inventory of trade secret information that is periodically updated
2. A security program to protect the technology at issue, often on a need-to-know basis with clear marking of information as "confidential, access restricted"
3. A written description of the security program that is provided to all employees

4. An enforcement officer or oversight procedure
5. An enforcement program, including litigation, if necessary, to enjoin unauthorized access or distribution

In the field of computing, these principles often mean that source code or other readable formats should be secured in a locked file and marked CONFIDENTIAL. All representations of the code as stored on magnetic or other media should be marked CONFIDENTIAL and secured. Computerized information should be password protected with restrictions on circulation of the password and periodic password changes. A notice of confidentiality should be displayed as soon as access to the program is obtained, with appropriate warnings on limitation of use. Levels of access should be controlled so that privileges to copy, read, and write are appropriately restricted. Surveillance of entries and logon should be routinely conducted to verify that there has been no unauthorized entry. Finally, periodic audits should be conducted to test and substantiate the security procedures.

For many years, each state developed its own brand of trade secret protection through evolving judicial decisions that establish something in this country called the common law, as distinguished from legislative enactments of a statute addressing the same issue. In 1985, the Uniform Trade Secrets Act (UTSA) was promulgated by the National Conference of Commissioners on Uniform State Laws, with one of its purposes to make uniform the rights and remedies available to a holder of a trade secret. This model law, however, needed to be adopted by each state before it became the law of the state. As of this writing, it has been adopted to some degree in 46 states with the exception of Massachusetts, New Jersey, New York, and Texas.

The UTSA defines a trade secret as information, including a formula, pattern, compilation, program device, method, technique, or process, that: (a) derives independent economic value, actual or potential, from *not being generally known to, and not being readily ascertainable by proper means by*, other persons who can obtain economic value from its disclosure or use; and (b) is the subject of efforts that are reasonable under the circumstances to maintain its secrecy. It also defines the unlawful taking of a trade secret, or misappropriation, as the wrongful use of a trade secret, including (a) knowingly acquiring the secret through improper means or (b) disclosing the secret without consent.

11.3.1 Remedies for Trade Secret Misappropriation. Misappropriation of a trade secret is the unauthorized use or disclosure of the trade secret. In simple parlance, it is a taking or theft. The taking can be by one who owes a fiduciary duty of confidentiality, such as an employee; it can be in breach of an agreement of confidentiality; or the taking can occur through improper access or means. The misappropriation can be treated under common law as the tort of conversion, trespass, unfair competition, or interference with contractual relations. As discussed, there are now specific statutory provisions under the UTSA for trade secret misappropriation. The UTSA grants the wronged party certain remedies that include enjoining the use of the misappropriated property, damages, and attorney's fees. When the misappropriation is of a physical item, such as a disk drive, the owner may ask the court to order seizure and return of its property.[10] In addition, where the misappropriation also violates other laws protecting intellectual property, such as where the taking infringes a copyright, the property owner may be entitled to additional relief.

Exactly what remedies are available will vary among the states. Interestingly, the very uniformity that the UTSA was intended to create has led to different treatment of

available claims and remedies. For example, before the UTSA, an employee's theft of the employer's confidential customer lists triggered a common law claim for breach of the implied fiduciary obligation owed by an employee to the employer as well as a claim for misappropriation of trade secrets. The UTSA provides that its remedies preempt other common law remedies; in other words, a claim under the UTSA trumps the claim for breach of fiduciary duty as well as the claim for misappropriation of trade secrets. There is a split in the courts as to whether the UTSA replaces only common law causes of action for misappropriation of trade secrets or extends to any tortious claims for relief that arise out of the misappropriation no matter how stated. The broader reach of the UTSA appears to be favored by the growing majority of courts that have considered this issue to date. The takeaway from this uncertainty is that computer security professionals should protect trade secrets, confidential information, and other valuable data through contractual terms with, among others, employees, vendors, and users to minimize the reliance on the UTSA.

In the event of a misappropriation, in addition to civil remedies, often separate state statutes treat the taking as a theft and a criminal act. Such statutes are generally state specific. Prior to 1996, the Trade Secrets Act (TSA) was the only federal statute prohibiting trade secret misappropriation. The TSA, however, was of limited utility because it did not apply to private sector employees and provided only limited criminal sanctions.[11] To combat an increase in computer crimes, Congress enacted the Economic Espionage Act of 1996 (EEA), which provided greater protection for the proprietary and economic information of both corporate and governmental entities against foreign and domestic theft.[12]

The EEA criminalizes two principal categories of corporate espionage: economic espionage and theft of trade secrets.[13] Section 1831 punishes those who steal trade secrets "to benefit a foreign government, foreign instrumentality, or foreign agent." Section 1832 is the general criminal trade secret provision.[14] The EEA criminalizes stealing, concealing, destruction, sketching, copying, transmitting, or receiving trade secrets without authorization, or with knowledge that the trade secrets have been misappropriated. It also criminalizes attempting to and conspiring to do any of these acts.[15] The EEA penalizes parties responsible for a taking that is intended to benefit a foreign government with fines up to $250,000 and imprisonment up to 10 years.[16]

The EEA explicitly defines a trade secret to include information stored in electronic media and includes "programs or codes, whether tangible or intangible" so long as:

> (a) the owner thereof has taken reasonable measures to keep such information secret; and (b) the information derives independent economic value, actual or potential, from not being generally known to, and not being readily ascertainable through proper means by, the public.[17]

Although one might assume that this definition is relatively straightforward, not everything is as it appears. In a case of domestic trade secret theft, the Court of Appeals for the Seventh Circuit examined what the EEA means when it says that the data or material "derives independent economic value, actual or potential, from not being generally known to, and not being readily ascertainable through proper means by, the public."[18] Noting that others had assumed that the word "public" meant the general public, the court in *Lange* astutely observed that this was not, in fact, the case. Moreover, the standard for measuring the persons who might readily ascertain the economic value of (in this case) the design and composition of airplane brake assemblies is not the average person in the street, for this assumes (as the court mentions) that any person can understand and apply something as arcane as Avogadro's number. Instead, the definition

of the term "the public" should take into account the segment of the population that would be interested in and understand the nature of that which has allegedly been misappropriated.

The international reach of the act is limited, extending outside of the United States only if: "(1) the offender is a natural person who is a citizen or permanent resident alien of the United States, or an organization organized under the laws of the United States or a State or political subdivision … or (2) an act in furtherance of the offense was committed in the United States."[19] Few defendants have been charged under the act since its passage in 1996, so the precise reach has yet to be tested. However, the language of the EEA applies its provisions to corporations with headquarters or operations subject to U.S. jurisdiction that could be prosecuted under the act. Finally, the remedies under the EEA can be invoked only by the United States. There is no private right of action under the act.

11.3.2 Vigilance Is a Best Practice. The key points of practice to remember are: Security and trade secret law are forever linked together. A trade secret cannot exist without such security. The maxim "Eternal vigilance is the price of liberty," often attributed to Thomas Jefferson, should in the context of business information protection be restated as "Eternal vigilance is the price of trade secret protection." It is not as catchy a phrase, but it is the price each business must pay if it relies in whole or in part on trade secret law for protection. In such situations, the greatest assurance of protection can be obtained through rigorous contractual terms and strenuous enforcement.

11.4 COPYRIGHT LAW AND SOFTWARE. Because of anxiety over the true extent of protection afforded software under patent and copyright law, software programs initially were protected as trade secrets. Such protection became increasingly problematic in today's society, where information technology and pressure for the free flow of information makes confidentiality controls more difficult to police. Copyright law now has evolved to include computer programs.

Since 1964, the United States Copyright Office has permitted registration of computer programs, although judicial decisions were divided on the applicability of the Copyright Act. In 1976, Congress passed the Copyright Act of 1976, which did little to resolve the ambiguity. Clarification finally was obtained in the Computer Software Copyright Act of 1980, which explicitly extended the protection of the copyright laws to software.[20] Any type of work that can be fixed in any tangible medium can be protected by copyright as literary works based on the authorship of the source and object code[21] even if the work can only be machine reproduced.

Copyright protection, however, does not protect "ideas."[22] Rather, it protects the particular expression of the idea. As can be seen by the parallel proliferation of spreadsheet programs, the idea for the spreadsheet program cannot be protected, but the particular code that produces the spreadsheet can be. In order to qualify for copyright protection, the work must be (a) original, (b) fixed in a tangible medium, and (c) not just the embodiment of an idea. Once obtained, copyright protection grants to the copyright owner the exclusive right to reproduce, to publish, to prepare derivative works, to distribute, to display, and to perform the copyrighted work. In 1990, Congress passed the Computer Software Rental Amendments Act,[23] which added to the list of copyright infringements the distribution of a computer program for commercial advantage. Materials copyrighted after 1978 are protected for the lesser of 75 years from the date of first publication or 100 years from the date of creation.

11.4.1 Works for Hire and Copyright Ownership. The copyright for a work does not always belong to the person who creates it. The most frequent exceptions are works that fall under the concept of a "work for hire." A work for hire is not owned by the creator but by the persons who hired the creator to create the work. Most often the concept applies to employees who have created a work *within the scope* of their employment. The key concept is the scope of employment. Even though a work is created outside of the office and normal working hours, it still will be a work for hire if it is within the scope of employment. However, a work that falls outside the scope of employment and that is created outside the office is likely not to be deemed a work for hire. Because of such issues, it is better practice when dealing with employees or independent contractors to provide specificity in an agreement as to what is a work and when the creation of a work will be governed by the doctrine of work for hire.

11.4.2 Copyright Rights Adhere from the Creation of the Work. Everyone who has looked at a copyrighted work is probably familiar with the symbol © affixed to any published copyrighted work, together with the name of the copyright holder and the year of creation or publication of the work. For many years, such notice was a fortiori necessary for copyright protection. Today, however, the copyright arises from the creation of a copyrighted work itself. It is still good practice to advise the world of potential infringement by inserting the formalities of a copyright on the work itself. In addition, one should register the work in the United States Copyright Office, which is currently developing a process for online registration. Registration of the copyright also permits one to claim statutory damages ranging from $500 to $20,000 for each violation, which often is useful to prevent additional infringements when no actual damages can be demonstrated. Moreover, in some jurisdictions, it may be necessary to register the copyright with the copyright office before one can actually sue to protect the copyright.

The change in copyright protection has interesting applications when applied to electronic works. The creation of the work in some permanent form is sufficient to trigger copyright protection. Thus, the creation of an electronic copy is sufficient permanency. What that means is that any electronic data are already conceivably subject to copyright protection at the time that they are viewed or received. Thus, in using any such information or "work," care must be taken that one does not infringe on a potential copyright without a license.

11.4.3 First Sale Limitation. The holder of the copyright has the right to sell or license the work. If the work is sold, the holder essentially loses all rights to control the resale of the work. This is known as the first sale doctrine. Once the item is placed in commerce, subsequent transfers cannot be restricted. The doctrine applies only to the copy that has actually been sold. It does not create a license to copy the item itself.

To avoid what sometimes can be a problem if the program winds up in the hands of a competitor, companies often prefer to license the item instead of selling it outright. If the work is licensed, only those rights that are contained in the license are transferred. All other rights of ownership remain with the licensor. Thus, a breach of the license gives the licensor of the copyrighted work the right to reclaim the work or prevent its further use or publication. However, if the license has all the basic indicia of a sale, it will be treated as one, notwithstanding the label.

One interesting intersection of these two principles is the requirement when upgrading software that the old version be present. As a condition of making the upgrade

available at a reduced rate, the seller normally requires that the older version be authenticated before the newer version can be installed. Such requirements are legal, as the owner of the earlier version could choose to sell it, but then would have to pay a higher price for the newer version and work to restrain subsequent sales of software from a user who expects to upgrade in the future.

11.4.4 Fair Use Exception. All copyright protection is subject to the doctrine of fair use.[24] Fair use permits the use of a work without authorization for a limited purpose. But what use constitutes fair use? The Copyright Act of 1976 suggested four, nonexclusive factors, for a court to consider:

1. What is the purpose and character of the use?
2. What is the nature of the copyrighted work?
3. How much of the copyrighted work is used?
4. What is the effect on the potential market for the work?

Despite its codification in the Copyright Act of 1976, fair use remains a nebulous doctrine—an equitable rule of reason, with each case to be decided on its own facts.[25] It is often misquoted and misapplied. The essential concept behind the doctrine of fair use is to permit public discussion, review, and debate of a copyrighted work without violating the copyright. Thus, the Copyright Act of 1976 gives as examples of fair use, situations of "criticism, comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research."

Fair use is not an antidote for failing to license a work. It should be invoked with care—understanding that the more material that is used and the more commercial the purpose, the less likely a court will find it applicable. Indeed, sometimes the only way to harmonize cases on whether a use is a fair use is to decide whether the court ultimately viewed the user as a "good" or "bad" guy.

11.4.5 Formulas Cannot Be Copyrighted. There are limitations on what expressions can be protected by copyright law. A frequent source of argument is whether, since one cannot protect the idea, the expression is directly driven by its content (i.e., the expression is simply a function of the idea). For that reason, formulas cannot be copyrighted.[26] This means that when formulas are part of a computer program, other modes of defense need to be considered, such as trade secret or possibly patent protection. If one were to disclose the formula through copyright publication, one would lose the ability to protect that information.

11.4.6 Copyright Does Not Protect the "Look and Feel" for Software Products. Copyright protection ordinarily extends to the physical manifestation of the computer program in the source code and object code. The operation of that code, as it translates to what the human mind perceives, has been described as the "look and feel" of the program. In attempting to quantify the concept of "look and feel," courts have considered whether the organization, structure, and sequence of the program can be protected. In the United States, *Whelan Associates, Inc. v. Jaslow Dental Lab., Inc.*,[27] gave the greatest extension to protecting look and feel. In that case, none of the code had been copied and the program operated on a different platform. Nonetheless, copyright infringement was found because the organization, structure, and sequence of the program had been copied. The court recognized that the structure and logic of

the program are the most difficult to create and that the idea could be protected as it was embodied in the program structure since, given the variety that was possible, the structure was not necessarily just an extension of the idea. Since *Whelan*, courts in the United States have retreated from such broad protection. In 1992, *Computer Associates, Inc. v. Altai, Inc.*[28] developed the so-called abstraction-filtration test. The results of that test define as unprotectable: (a) program structures that are dictated by operating efficiency or functional demands of the program and therefore deemed part of the idea and (b) all tools and subroutines that may be deemed included in the public domain. Only what remains is to be compared for possible copyright infringement.

While protection of look and feel may vary among the different federal circuits, in general, the courts are swinging away from broader protection. However, this may not necessarily be true internationally; English law appears to grant the broader protections afforded by the *Whelan* decision.

11.4.7 Reverse Engineering as a Copyright Exception. Within the field of computer software, cases have considered whether "dissection" in order to reverse engineer the program is a violation of the copyright. To those involved in protecting software programs, as well as those involved with competing products, the answer appears to be that reverse engineering does not constitute an infringement, even though the disassembly of the program falls squarely within the category of acts prohibited by the Copyright Act because of the doctrine of fair use. The Ninth Circuit in *Sega Enterprises Ltd. v. Accolade, Inc.*[29] found as a matter of law that:

> … where disassembly is the only way to gain access to the ideas and functional elements embodied in a copyrighted computer program and where there is a legitimate reason for seeking such access, disassembly is a fair use of the copyrighted work.[30]

The Ninth Circuit is not the only circuit that has upheld reverse engineering against a copyright claim. The Federal Circuit reached a similar conclusion regarding reverse engineering of object code to discern the "ideas" behind the program in *Atari Games Corp. v. Nintendo of America, Inc.*[31] The fair use rationale of *Sega* was also adopted by the Eleventh Circuit in *Bateman v. Mnemonics, Inc.*[32] on the grounds that it advanced the sciences. In addition, in *Assessment Techs. of WI, LLC, v. WIREData, Inc.*, the Seventh Circuit relied on *Sega* and determined that WIREData, Inc. could extract uncopyrighted data from a copyrighted computer program, noting that the purpose of the extraction was to get the raw data, not compete with Assessment Technologies by selling copies of the program itself.[33] In *Evolution, Inc. v. SunTrust Bank*, the Tenth Circuit relied on both *Sega* and *WIREData* when it allowed the defendant to copy part of plaintiff's source code to extract uncopyrighted data from plaintiff's copyrighted computer program.[34] Thus, unless careful thought is given to the application of copyright protection, merely copyrighting the software will not necessarily protect against imitation.

11.4.8 Interfaces. There is an open issue as to whether copyright protects the format for interfacing between application and data. Competitors, particularly in the area of gaming, look to reverse engineer the interface format to make new modules compatible with existing hardware. Such reverse engineering has been held not to violate the copyright laws, so long as the new product does not display copyrighted images or other copyrightable expressions.[35] Thus, the nonprotectable interface may be protected if such copyrighted images or expressions are embedded in the display.

11.4.9 Transformative Uses. One of the factors that the doctrine of fair use considers is the "amount and substantiality of the portion used in relation to the copyrighted work as a whole."[36] In practical terms, this means that courts look at how much was taken and for what purpose. One could take a little but still take the essence of the program. One could also take a little that did not attempt to duplicate but rather used the copyrighted material as a springboard for a new creation. Out of this qualitative and quantitative investigation comes the notion of transformative use, which became the coin of analysis in the Supreme Court's 1994 decision in *Campbell v. Acuff-Rose Music, Inc.*[37] *Campbell* addressed the concept in terms of a claim of copyright infringement involving a rap parody of a popular song. There, taking its clues from the opening language of Section 107 codifying fair use, the Supreme Court asked whether the "new" work "adds something new, with a further purpose or different character, altering the first with new expression, meaning or message; it asks, in other words, whether and to what extent the new work is transformative."[38] The Court then laid down the test to be applied.

> Although such transformative use is not absolutely necessary for a finding of fair use, … the goal of copyright, to promote science and the arts, is generally furthered by the creation of transformative works. Such works thus lie at the heart of the fair use doctrine's guarantee of breathing space within the confines of copyright, … and the more transformative the new work, the less will be the significance of other factors, like commercialism, that may weigh against a finding of fair use.[39]

Thus, a transformative use may play off of a prior copyright and still not be deemed an infringement so long as the resulting new work is just that—new.

11.4.10 Derivative Works. Under Section 106(2) of the Copyright Act of 1976, the copyright owner has the exclusive right "to prepare derivative works based upon the copyrighted work." The act defines a "derivative work" as:

> … a work based upon one or more pre-existing works, such as a translation, musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgement, condensation, or any other form in which a work may be recast, transformed, or adapted. A work consisting of editorial revisions, annotations, elaborations, or other modifications which, as a whole, represent an original work of authorship, is a "derivative work."

A "derivative work" is thus defined as an original work that is independently copyrightable. To infringe the exclusive right to prepare a derivative work granted by the Copyright Act to the copyright owner, the infringer need not actually have copied the original work or even have fixed in a tangible medium of expression the allegedly infringing work.[40] The right, therefore, to create the derivative work can be a useful tool in counterbalancing attempts to pirate computer programs and the issue of fair use.

The Copyright Act creates an exemption for a lawful owner of a purchased license for a computer program to adapt the copyrighted program if the actual adaptation "is created as an essential step in the utilization of the computer program in conjunction with a machine and it is used in no other manner."[41] The adaptation cannot be transferred to a third party. The right to adapt is, in essence, the right to modify or, in the language of the act, to create a derivative work. Such changes can be made even without the consent of the software owner so long as such modifications are used only internally and are necessary to the continuing use of the software.[42]

11.4.11 Semiconductor Chip Protection Act of 1984. The Semiconductor Chip Protection Act of 1984 (SCPA) protects as part of the Copyright Act "mask works fixed in a semiconductor product."[43] The SCPA protects not the product itself but the copying of the circuit design or blueprint. Because of reverse engineering, the protections afforded by SCPA are limited in practice.

11.4.12 Direct, Contributory, or Vicarious Infringement. Copyright infringement generally requires a showing of substantial similarity between allegedly offending use and the protected expression contained in a work. Infringement can occur through the simple act of printing (without permission), by posting on the Web or other form of unauthorized distribution, by creating a derivative work, or by another act that interferes with the copyright holder's rights.

A copyright can be infringed directly, contributorially, or vicariously. Direct infringement is the term ascribed to the actor who violates the copyright. Contributory infringement involves knowingly providing the means for the violation to occur. Liability for contributory infringement may be predicated on actively encouraging (or inducing) infringement through specific acts, or on distributing a product that distributees use to infringe copyrights, if the product is not capable of "substantial" or "commercially significant" noninfringing uses.[44] But secondary liability for copyright infringement does not exist in the absence of direct infringement by a third party. Vicarious infringement occurs when one is responsible for or controls the actions of another who violates the infringement. The usual situation is that of an employer's responsibility for the acts of an employee.

Not all situations admit themselves of simple answers, as when a person commits direct infringement by actually photocopying a work. New technologies constantly pose issues as to whether infringement has occurred and whether the infringement violates the public interest. In general, when faced with an issue of potential copyright infringement, the questions to ask are:

- Can the product or service be used to infringe a copyright, or is the product capable of substantial noninfringing uses?
- If so, did the owner of the product or service encourage the user to use it for infringement?
- Alternatively, did the owner of the product or service have knowledge of the specific infringing use and have the ability to prevent it?

Today, we take Internet service providers (ISPs) for granted. But application of these questions initially led courts to conclude that ISPs were liable for contributory infringement. For example, a Web site that encouraged and facilitated the uploading of copyrighted materials was found to be a direct infringer of the copyright of the owner even though the provider did not actually do the uploading.[45] Similarly, an ISP that was notified of a copyright violation that was posted on its server and failed to correct it could be found to have contributory liability for the infringement.[46] In its wisdom, however, Congress, in the Digital Millennium Copyright Act (DMCA), created a safe harbor for ISPs so that, as a matter of public policy, an ISP does not have to monitor each and every transmission for potential copyright infringement.

11.4.13 Civil and Criminal Remedies. The Copyright Act contains several sections that specifically address the penalties and remedies for infringement. They

include injunctive relief (i.e., a court order terminating the infringing conduct),[47] impounding and disposing of infringing articles,[48] damages,[49] litigation costs and attorneys' fees,[50] and criminal penalties.[51] Although this chapter cannot address all of the permutations of remedies and penalties available, a few are worth mentioning.

Generically, a copyright owner must choose between its actual losses (i.e., what it actually lost and any profits realized by the infringer) and statutory damages.[52] Actual damages imply economic losses actually suffered as a result of the infringement. The kinds of actual damages that have been awarded include development costs of the software,[53] the economic consequences of lost customers,[54] lost future sales,[55] the value of the infringer's licensing fees where the licensor is precluded from market sales,[56] lost market value of the material infringed,[57] and lost royalty payments.[58] An award of actual damages is not automatic; the license holder has the burden of proving that the infringing activity and the economic loss are causally connected, at which point the infringing party must show that the license holder would have incurred the loss anyway.[59]

A copyright owner may elect to receive statutory damages rather than actual damages and the infringer's profits.[60] Making the election is mandatory, and it must be done before final judgment is entered. Once the election is made, it is final. The statutory damages generally range from $500 to $20,000 "for all infringements involved in the action, with respect to any one work, for which any two or more infringers are liable jointly and severally. … For purposes of this section, all the parts of a compilation or derivative work constitute one work."[61] This amount may be increased to $100,000 if the court finds that the infringement was willful and reduced to $200 if the court finds that the infringer "was not aware and had no reason to believe" that the act was an infringement.[62]

Statutory damages theoretically[63] are intended to approximate the actual damages suffered, and were crafted as an alternative compensation scheme for copyright owners, when actual damages are difficult to calculate. In determining whether to elect actual or statutory damages, a copyright owner ought to perform a careful analysis to determine how many separate infringements occurred that justify, under the statute, separate awards. Although posting different copyrighted computer software programs on a bulletin board for downloading constitutes multiple infringements,[64] making multiple copies of the same cartoon character in different poses constitutes a single infringement because only one work was copied.[65]

As mentioned, this is one of the statutory schemes that discourage frivolous litigation by imposing the cost of litigating on the losing party. The statute permits the substantially prevailing party to recover its reasonable attorneys' fees and costs from the losing party. Who is the substantially prevailing party and what constitutes reasonable attorneys' fees are separate and distinct issues that will be decided by the courts.

Copyright violations also can be criminally prosecuted, and generally require demonstration of *mens rea,* or intent. One or more infringements having a total retail value of more than $1,000 within a 180-day period or "for purposes of commercial advantage or private financial gain" can be punished by one to five years of imprisonment and fines. Even without demonstration of a motive of financial gain, 10 or more infringements having a value in excess of $2,500 can result in up to three years in jail and fines. Repeated violations carry stiffer penalties. Finally, one who knowingly aids or abets a copyright infringement is also subject to criminal prosecution.

11.5 DIGITAL MILLENNIUM COPYRIGHT ACT. In 1998, Congress passed the Digital Millennium Copyright Act (DMCA) to address concerns raised by the

Internet and copyright issues in the context of our increasingly technological society. The DMCA creates a civil remedy for its violation as well as criminal penalties starting after October 2000. One of the purposes of the DMCA is to protect the integrity of copyright information. Removal of a copyright notice, or distribution knowing that such copyright has been removed, is now actionable.[66]

11.6 CIRCUMVENTING TECHNOLOGY MEASURES. Article 11 of the World Intellectual Property Organization Copyright Treaty required all signatory countries to provide adequate legal protection and remedies against the circumvention of technical measures intended to secure copyrights. In response, Congress adopted Section 1201 of the DMCA, which generally prohibits the act of circumventing, and trafficking in the technology that enables circumvention of, protection measures designed to control access to copyrighted work.[67] Both civil and criminal remedies also now exist under the DMCA if one circumvents "a technological measure that effectively controls access to a work protected" by the Copyright Act.[68] It is a civil violation and a crime to "manufacture, import, offer to the public, provide or otherwise traffic in any technology, product, service, device, component, or part thereof" that "is primarily designed or produced for the purpose of circumventing a technological measure that effectively controls access to a work protected" under the Copyright Act.[69] A technological measure effectively controls access to a work if the measure, "in the ordinary course of its operation, requires the application of information or a process or a treatment, with the authority of the copyright owner, to gain access to the work."[70] One circumvents such technology measure if one uses a means "to descramble a scrambled work, to decrypt an encrypted work, or otherwise to avoid, bypass, remove, deactivate, or impair a technological measure," without the authority of the copyright owner.[71]

In *RealNetworks, Inc. v. Streambox, Inc.*,[72] Streambox distributed software that enabled users to bypass the authentication process employed by RealNetworks, which distributes audio and video content over the Internet. Thus, Streambox users could get the benefit of the RealNetworks streaming audio and video content without compensating the copyright owners. The United States District Court in Washington State found that the Streambox software was a technological measure that was designed to circumvent the access and copy control measures intended to protect the copyright owners.[73]

In a case involving digital video disc (DVD) encryption, a U.S. District Court in New York enjoined posting links to sites where visitors may download the decryption program as trafficking in circumvention technology and a violation of the DMCA.[74] In *Universal City Studios, Inc. v. Reimerdes*, the court rejected an argument that the use of the decryption software constituted free expression protected by the First Amendment of the U.S. Constitution. On appeal, the appellant argued that the injunction violated the First Amendment because computer code was speech, was entitled to full protection, and was unable to survive the strict scrutiny given to protected speech.[75] The appellate court found that the computer code used in the program was protected speech:

> Communication does not lose constitutional protection as "speech" simply because it is expressed in the language of computer code. Mathematical formulae and musical scores are written in "code," i.e., symbolic notations not comprehensible to the uninitiated, and yet both are covered by the First Amendment. If someone chose to write a novel entirely in computer object code by using strings of 1's and 0's for each letter of each word, the resulting work would be no different for constitutional purposes than if it had been written in English. The "object code" version would be incomprehensible to readers outside the programming community (and tedious to read even for most within the community), but it would be no more incomprehensible

> than a work written in Sanskrit for those unversed in that language. The undisputed evidence reveals that even pure object code can be, and often is, read and understood by experienced programmers. And source code (in any of its various levels of complexity) can be read by many more. See *Universal I*, 111 F. Supp. 2d at 326. Ultimately, however, the ease with which a work is comprehended is irrelevant to the constitutional inquiry. If computer code is distinguishable from conventional speech for First Amendment purposes, it is not because it is written in an obscure language.[76]

The court then analyzed the type of scrutiny that should be applied where the restriction is content neutral:

> Having concluded that computer code conveying information is "speech" within the meaning of the First Amendment, we next consider, to a limited extent, the scope of the protection that code enjoys. As the District Court recognized, *Universal I*, 111 F. Supp. 2d at 327, the scope of protection for speech generally depends on whether the restriction is imposed because of the content of the speech. Content-based restrictions are permissible only if they serve compelling state interests and do so by the least restrictive means available. *See Sable Communications of California, Inc. v. FCC*, 492 U.S. 115, 126, 106 L. Ed. 2d 93, 109 S. Ct. 2829 (1989). A content-neutral restriction is permissible if it serves a substantial governmental interest, the interest is unrelated to the suppression of free expression, and the regulation is narrowly tailored, which "in this context requires … that the means chosen do not 'burden substantially more speech than is necessary to further the government's legitimate interests.'" *Turner Broadcasting System, Inc. v. FCC,* 512 U.S. 622, 662, 129 L. Ed. 2d 497, 114 S. Ct. 2445 (1994) (quoting *Ward v. Rock Against Racism*, 491 U.S. 781, 799, 105 L. Ed. 2d 661, 109 S. Ct. 2746 (1989)).[77]

Finding that the government's interest in preventing unauthorized access to encrypted copyrighted material is unquestionably substantial, and that the regulation of decryption programs served that interest, the appellate court upheld the prohibitions against both posting of, and linking to, the decryption program.

Not all efforts to "circumvent" restrictions, however, come within the prohibitions of the DCMA. In *I.M.S. Inquiry Mgmt. Sys. v. Berkshire Info. Sys.,*[78] the defendant had used a valid password provided to plaintiff's own customers and user identification to view plaintiff's e-Basket system exactly as the customer itself might have done. The court concluded that although this might be viewed as a technology measure, it was not circumvention of a digital wall within the meaning of the DCMA.

11.6.1 Exceptions to the Prohibitions on Technology Circumvention.

The DMCA, however, explicitly carves out all defenses to copyright infringement, including the doctrine of fair use, as being unaffected by the passage of the DMCA. In some circumstances fair use can include reverse engineering.

11.6.1.1 Fair Use and Reverse Engineering.

Thus, one can spy through reverse engineering still without running afoul of copyright protection or the DMCA.

However, in *Bowers v. Baystate Technologies, Inc*.,[79] a split Federal Circuit Court of Appeals found that a shrink-wrap license prohibiting reverse engineering was enforceable against the licensee who had reverse engineered Bowers's CAD Designer's Toolkit to develop a competing product. The *Bowers* court found that the contractual language trumped the "fair use" permitted under the Copyright Act. The Fifth Circuit has reached the opposite result in the earlier decision of *Vault Corp. v. Quaid Software, Ltd*.,[80] specifically finding that the Copyright Act preempts state law that attempts to prohibit disassembly, and holding a mass distribution license agreement unenforceable.

Thus, the extent to which *Bowers* may be followed is still unclear, but it appears to be questioned in subsequent decisions.[81] *Bowers* suggests a course that businesses can attempt to follow to curtail reverse engineering, which is to limit that right by contract. If *Bowers* becomes widely accepted, the United States will be in conflict with the European Union on this issue. In its 1991 Software Directive, the European Union set forth a right to reverse engineer that is consonant with "fair use" under the Copyright Act. The Software Directive also provided that the right cannot be waived by contract. So, until *Bowers* is settled, if a shrink-wrap license prohibits reverse engineering, it would be best to consider engaging in such activity abroad.

11.6.1.2 Other Exceptions. The DMCA also creates an important exception that recognizes the right to reverse engineer if (a) the person has lawfully obtained the right to use a copy of a computer program, and (b) the sole purpose of circumventing the technology measure is to identify and analyze "those elements of the program that are necessary to achieve interoperability of an independently created computer program with other programs."[82] The DMCA creates a similar exemption for circumvention for the purpose of "enabling the interoperability of an independently created computer program with other programs, if such means are necessary to achieve such interoperability."[83] The term "interoperability" is defined to encompass the "ability of computer programs to exchange information and of such programs mutually to use the information which has been exchanged."[84] The information acquired through these permitted acts of circumvention may also be provided to third parties so long as it is solely used for the same purposes.[85]

Circumvention is permissible under these exemptions, however, "only to the extent [that it] does not constitute copyright infringement." Two cases, *Chamberlain Group, Inc. v. Skylink Techs., Inc.*,[86] and *Lexmark Int'l, Inc. v. Static Control Components, Inc.*,[87] are particularly instructive. In both cases, the courts permitted a competitor's access and reverse engineering under this exemption. In contrast, in *Storage Tech Corp. v. Custom Hardware Engineering Consulting, Inc.* (D. Mass. 2004), the defendant bypassed a protective access key to activate the diagnostics program by copying the code into the random access memory (RAM) of the defendant's access device. The District Court found that this copying constituted infringement. The result was reversed in a 2 to 1 decision in the United States Federal Circuit[88] based on a reading of sections 117(a) and (c) of the DMCA, which permits copying for maintenance purposes. This string of decisions has led to recommendations that access be controlled by a method that would cause copyright infringement and that access protect not just the copyrighted program but copyrighted data so as to exclude the rationale of the Federal Circuit. Suggestions have been made that certain parts of copyrighted executable code be encrypted and that a decryption key be required that will create a copy of the code and protected data as part of the process so as to create an argument of copyright infringement. These types of recommendations remain untested, and the simpler course may be control through terms inserted into the licensing agreement.

Exempt from the DMCA, as well, are "good faith" acts of circumvention where the purpose is encryption research. A permissible act of encryption research requires that (a) the person lawfully have obtained a copy, (b) the act is necessary to the research, (c) there was a good faith effort to obtain authorization before the circumvention, and (d) such act does not constitute an infringement under a different section of the Copyright Act or under the Computer Fraud and Abuse Act of 1986. With the caveat that it must be an act of good faith encryption research, the technological means for circumvention can be provided to others who are working collaboratively on such

research. The issue of good faith encryption research looks to what happened to the information derived from the research. If it was disseminated in a manner that was likely to assist infringement, as opposed to reasonably calculated to advance the development of encryption technology, then the act still falls outside of the exemption. Other factors that go into the determination of good faith are whether the person conducting the research is trained, experienced, or engaged in the field of encryption research and whether the researcher provides the copyright owner with a copy of the findings.

The DMCA also has a bias against the collection or dissemination of personally identifying information. Thus, it is not a violation of the DMCA to circumvent a technology measure that essentially protects, collects, or disseminates personally identifying information, provided that the circumvention has no other effect, and provided that the program itself does not contain a conspicuous notice warning against the collection of such information and a means to prevent or restrict such collection.[89] In short, one can disable cookies if the program does not itself permit a user to do so.

Finally, insofar as relevant to this chapter, the DMCA also excludes from its scope "security testing." The DMCA grants permission to engage in security testing that, but for that permission, would violate the terms of the DMCA. If the security testing, for some reason, violated some other provision of the Copyright Act or the Computer Fraud and Abuse Act of 1986, then it is still an act of infringement. The DMCA, in part, considers whether a violation occurred, and by whom the information was used. The factors to be considered include if the information was used to promote the security of the owner or operator of the computer network or system, if it was shared with the developer, and if it was used in a manner that would not facilitate infringement.[90] For purposes of the DMCA, security testing means accessing either an individual computer or network for the purpose of "good faith testing, investigating, or correcting, a security flaw or vulnerability, with the authorization of the owner or operator."[91]

11.6.1.3 Remedies. The criminal penalties for violation of the DMCA can be quite severe. If the violation is willful for commercial gain, the first offense bears a fine of up to $500,000 or 5 years' imprisonment. Subsequent violations bear fines of up to $1 million dollars or 10 years imprisonment. Civil remedies include an order to restrain the violation, damages for lost profits, damages for recovery of the infringer's profits, or statutory damages for each violation. Depending on the section of the DMCA at issue, each violation can generate fines of up to $2,500 or $25,000. Since each act of infringement can constitute a violation, the statutory fines can become quite substantial.

11.7 PATENT PROTECTION.

Ideas, which are not protected by copyright, can be protected through a patent. In general, the patent laws protect the functionality of a product or process.

11.7.1 Patent Protection Requires Disclosure.

A patent can be properly obtained if the invention is new, useful, nonobvious, and disclosed. The patent exchanges a grant of an exclusive monopoly over the invention in return for disclosure. Disclosure is the trigger point for patentability. The disclosure supports the claims of patentability (i.e., it sets up the claim that the invention is both new and nonobvious) and also the scope of what can be protected. Thus, 35 U.S.C. section 112 provides:

> The specification shall contain a *written description* of the invention, and of the manner and process of making and using it, in such full, clear, concise and exact terms as to *enable any person skilled in the art* to which it pertains, or with which it is most nearly connected, to make

> and use the same, and shall set forth the *best mode* contemplated by the inventor of carrying out his invention.
>
> The specification shall conclude with one or more claims particularly pointing out and *distinctly claiming* the subject matter which the applicant regards as his invention. [Emphasis added.]

A patent therefore must disclose the best mode for implementing the invention, a clear written description of the invention, sufficient detail so that a practitioner can understand and make use of the description, and distinct claims, in order for a patent to issue.[92] Through adequate disclosure of the invention, the application gives notice of the technology involved in the patent so as to put the public on fair notice of what would constitute an infringement. From a public policy perspective, the disclosure enlarges the public knowledge. From the inventor's perspective, the trade-off is disclosure for exclusivity. Depending on how the invention is to be used and the areas in which protection will be necessary, disclosure may not be the best means of protecting the invention. This is particularly true if the inventor is not convinced it will be deemed nonobvious from prior art, in which case it will be subject to challenge, or if, after disclosure, other companies may legally use the disclosed information for competitive advantage. The effects of disclosure should be carefully considered before applying for patent protection.

11.7.2 Patent Protection in Other Jurisdictions. Patent protection is jurisdictional. What that means, in general, is that a patent has legal meaning in the country that granted it. The United States is a signatory to the Paris Convention for the Protection of Industrial Properties, which has roughly 160 signatories. The Paris Convention essentially grants a one-year grace period for filing national patent applications in each selected signatory, to obtain the benefit of the original filing date in the United States. An alternative, open to members of the Paris Convention, is the Patent Cooperation Act. This permits the filing of an international patent that basically gives the patentee an 8- to 18-month window to test feasibility, and which simplifies the national application process.

11.7.3 Patent Infringement. Like the remedies for copyright infringement, the remedies for patent infringement include injunctive relief and damages that, by statute, are not less than a reasonable royalty for the infringing use.[93] If the infringement is willful, the damages can be trebled. Attorneys' fees can be awarded, but only in exceptional cases.

In the area of exported computer software, an issue of note has arisen under 35 U.S.C. section 271(f). Section 271(f) was added in 1984 to the patent law to prevent infringers from avoiding liability by finishing goods outside of the United States. An infringer will be liable if its intent is to manufacture or supply a component from the United States to be combined elsewhere, if it would be an infringement had it occurred within the United States. Exported software may be considered a "component" under section 271(f). In *Microsoft Corp. v. AT&T Corp.*,[94] the issue was whether a master disk supplied by Microsoft abroad for duplication and installation abroad of its Windows program ran afoul of AT&T's patent. In overruling the Federal Circuit, the Supreme Court concluded that it did not.

> Section 271(f) prohibits the supply of components "from the United States … in such manner as to actively induce the combination of such components." § 271(f) (1). Under this formulation, the very components supplied from the United States, and not copies thereof, trigger § 271(f)

> liability when combined abroad to form the patented invention at issue. Here, as we have noted, the copies of Windows actually installed on the foreign computers were not themselves supplied from the United States. Indeed, those copies did not exist until they were generated by third parties outside the United States. Copying software abroad, all might agree, is indeed easy and inexpensive. But the same could be said of other items: "Keys or machine parts might be copied from a master; chemical or biological substances might be created by reproduction; and paper products might be made by electronic copying and printing."… The absence of anything addressing copying in the statutory text weighs against a judicial determination that replication abroad of a master dispatched from the United States "supplies" the foreign-made copies from the United States within the intendment of § 271(f).

Unless section 271(f) is amended, it may have profound implications for subverting the ability of a U.S. company to control patent infringement where software is a component of a patented invention.

11.8 PIRACY AND OTHER INTRUSIONS. For as long as ideas and innovation have been a source of commercial or social value, the terms on which these ideas and innovations have been available for use and exchange by others has been the subject of significant tension. Although inventors and creators of commercially viable products and processes want to maximize the return on their investment, marketplace pressure for cost efficiency (often motivated by human and corporate greed) fuels a constant drive to remove the inventors' and creators' royalties from the cost of production. Thus, the ancient notion of piracy, the unauthorized boarding of a ship to commit theft, and the unauthorized use of another's invention or production[95] remains alive and well. The piracy we speak of is not simply the unauthorized copying of millions of compact discs (CDs); increasingly it includes the unauthorized scraping of data from Web sites, abuse of authorized Internet use, theft of employee data, and similar activities.

11.8.1 Marketplace. The demand for unlicensed access to and use of software and entertainment media increases annually. In its 2007 survey regarding computer security among corporate and governmental institutions, the Computer Security Institute and the U.S. Federal Bureau of Investigation found that 59 percent of all respondents discovered employees who abused Internet privileges for a variety of unauthorized purposes.[96] A 2007 study by the Software & Industry Information Association reported worldwide revenue loss from the piracy (unlawful copying and distribution) of software exceeding $28.8 billion in 2007.[97] In countries such as China, despite recent overtures to the contrary, piracy is not merely sanctioned, it constitutes an investment by government agencies.[98]

In recent years, in large part due to the saturation of Internet access, there has been a tremendous proliferation of technologies designed to access and distribute (without authorization) protected software applications and entertainment media. This has posed a tremendous challenge for license holders, legislators, and law enforcement authorities. The results have included attempts to punish both unauthorized access and use of protected material. In the process, there has been a transformation in the definition of what is protected and some confusion about the extent of that protection when the Internet is involved.

11.8.2 Database Protection. Databases, the organized compilation of information in an electronic format, are prominent elements of any discussion concerning copyright protection. Compilations of information, data, and works are protectable under the Copyright Act.[99] To secure copyright protection for a compilation, a party

must demonstrate that (1) it owned a valid copyright in the compilation; (2) the alleged infringer copied at least a portion of the compilation; and (3) the portion so copied was protected under the Copyright Act.[100] In this context, the Copyright Act protects the "original" selection, coordination, or arrangement of the data contained in the compilation.[101]

To the extent that compilations contain purely factual information (e.g., existing prices of products and services), there is no protection because the facts themselves lack originality.[102] It does not matter that the author "created" the facts of the prices being charged for the product or service.[103] To sustain a claim of copyright protection for compilations of fact, the author must demonstrate creativity in the arrangement of the data. Standard or routine arrangements are likewise beyond the act's umbrella.[104] This is in contrast to the European Union's Database Directive, which does not require creativity as an element for the protection of a database. Rather it protects investment in databases under copyright protection subject, however, to fair use qualifications.

The United States Supreme Court has held that the compilation into a database of original works by contributing authors to newspapers and magazines violates the copyrights of the individual authors when the database does not reproduce the authors' articles as part of the original collective work to which the articles were contributed. In *New York Times Co., Inc. v. Tasini*,[105] authors who contributed articles and other works to the *New York Times*, *Time* magazine, and *Newsday* sued when they learned that the articles that they sold to the publishers for use in the respective publications were being reproduced and made available online, through LEXIS/NEXIS, an online database, and on CD-ROM. In most instances, the reproductions were of the individual articles outside of the newspaper or magazine context, in a collection of works separately protected by the Copyright Act. The Supreme Court held that, because the publishers of the new collective works made no original or creative contribution to the individual authors' original works, they could not reproduce and distribute those works outside of the format that each publisher created for the original collections of works, without permission from, or payments to, each author.[106]

11.8.3 Applications of Transformative and Fair Use. The concepts of transformative use and fair use (to the extent that they are separable) discussed earlier in this chapter have played a substantial role in recent decisions involving the authorized use of electronic media and the Internet. The starting point for this application of the doctrine is the U.S. Supreme Court's decision in *Sony Corporation v. Universal City Studios, Inc.*,[107] the famous battle over Betamax initiated by the movie industry. At issue was whether electronic recording machines could record television programs to permit individuals to "time-shift" television programs (i.e., to record programs for viewing at a time other than the time of airing). In its decision, the *Sony* Court found that time shifting was a productive use of the television programs for a purpose other than the original commercial broadcast, and was not an attempt either to duplicate the original purpose or to impact the commercial market for these programs. The Court emphasized the noncommercial element inherent in time shifting.[108]

11.8.4 Internet Hosting and File Distribution. The growth of the breadth and scope of the Internet has been accompanied by increasing questions about the extent to which the distribution of otherwise protected expressions change their form when converted into an electronic format. These questions arise for ISPs, which provide the pathway for distributing protected material, and for end users who post such materials

on their Web sites and bulletin boards. For ISPs, the DMCA provides some initial comfort.

Title II of the DMCA, designated the "Online Copyright Infringement Limitation Act," establishes several infringement liability safe harbors for service providers. The "Information residing on systems or networks at direction of users"[109] safe harbor is available to any provider of "online services or network access, or the operator of facilities thereof, . . ." including "digital online communications, between or among points specified by user, of material of the user's choosing, without modification to the content of the material as sent or received"[110] that "has adopted and reasonably implemented, and informs subscribers and account holders of the service provider's system or network of, a policy that provides for the termination in appropriate circumstances of subscribers and account holders of the service provider's system or network who are repeat infringers" and "accommodates and does not interfere with standard technical measures."[111] To qualify for the safe harbor, the service provider must demonstrate that:

1. It has no actual or constructive knowledge that information on its system is infringing, it is not aware of circumstances from which infringement is apparent or, upon obtaining such knowledge or awareness, it acts expeditiously to remove those materials;
2. It receives no financial benefit directly attributable to the infringing activity, *and*
3. Upon receipt of a notice of infringing material on its system, responds expeditiously to remove, or disable access to, the material.[112]

Assuming that the safe harbor does not apply (as, for instance, because the ISP failed to act on a notice of infringing activity), many service providers may nonetheless escape liability. In the first, and seminal, case on this topic, *Religious Technology Center v. Netcom On-Line Communication Services, Inc.*,[113] an ISP hosted a bulletin board service on which Church of Scientology publications were posted by a former minister. The District Court held that the ISP must demonstrate that its use was of public benefit (facilitating dissemination of creative works including, but not limited to, the infringing work); that its financial gain was unrelated to the infringing activity (e.g., subscription fees from providing email systems rather than fees from the display or sale of the infringing work); that its use was unrelated to the use of the owner of the work; that the ISP copied only what was necessary to provide its service; and that its use of the material had no demonstrable effect on the potential market for the work.[114] In *CoStar Group, Inc. v. LoopNet, Inc.,* the Fourth Circuit relied on *Netcom*, its codification in the DMCA, and the fact that the DMCA does not limit the application of other infringement defenses, and held that "the automatic copying, storage, and transmission of copyrighted materials, when instigated by others, does not render an ISP strictly liable for copyright infringement under §§501 and 106 of the Copyright Act."[115]

For Web site owners and users who post allegedly infringing material, the courts have had much less difficulty discarding the transformative fair use arguments. This has been particularly true in the purely commercial setting, as where the infringing party gains direct financial benefit from the infringing material,[116] and where the posted material is an exact copy of the protected work without any transformation to something creative or original.[117] In a case that goes to the heart of the open-access nature of the Internet, one court recently held that a copyright owner who posts its work on the Internet for

free distribution as shareware may defeat a transformative fair use defense by also posting an express reservation of distribution rights.[118]

11.8.5 Web Crawlers and Fair Use. The Internet, premised on open exchange of data and economic efficiency, has spawned a spate of data search and aggregation software tools that scan the Web looking for information requested by the user. The process used by these search engines[119] includes identifying data on the Web that conforms to the search parameters and then downloading that data. Since the copying usually occurs without the express permission of the copyright owner, some have argued that such copying constitutes an infringement. Although there is very little precedent concerning the application of transformative fair use to automated data retrieval systems, at least one court has upheld the use of the defense to an infringement claim.[120]

11.8.6 HyperLinking. In *Perfect 10 v. Google, Inc.*,[121] affirmed in part and remanded in part, *Perfect 10, Inc. v. Amazon.com, Inc*.,[122] Perfect 10 (P10) claimed that Google was infringing its ownership of copyrights in certain images and thumbnails hosted by third-party and P10's Web sites when Google's image search picked them up for display as framed full-size images and as thumbnails on computers and cell phones. The court concluded that hyperlinking did not constitute display for purposes of direct copyright infringement. On appeal, the case was remanded for further consideration as to whether the conduct fell within the general rule for contributory liability. To appreciate the context in which the courts are wrestling with these issues in the light of new technology, a review of the District Court's analysis should be studied.

11.8.7 File Sharing. Transformative fair use will not protect the verbatim retransmission of protected work in a different medium when there is a substantial and detrimental impact on the market for the protected work. In *A&M Records, Inc. v. Napster, Inc.*,[123] Napster enabled users to share music files over the Internet by downloading the file-sharing software to their hard drive, using the software to search for MP3 music files stored on other computers, and transferring copies of MP3 files from other computers. The court of appeals held that Napster users were merely retransmitting original works in a different medium and that this did not constitute a transformation of the original work. The court also found that sharing of music files over the Internet had, and would have, a significant and detrimental impact on the existing and potential market for CDs and digital downloads of the copyright owners' works. Picking up on the *Sony* decision's emphasis on the distinction between commercial and personal use, the Court of Appeals found that Napster's Web site effectively made the works available for use by the general public and not simply for the personal use of individual users.[124]

Napster's demise, however, did not end the controversy over file sharing. Trying to avoid Napster's method of directly enabling file sharing, entities such as Grokster and StreamCast developed software creating peer-to-peer networks through which individual computers communicate to exchange files without the necessity of a central server.[125] The Supreme Court recently revisited copyright infringement and file sharing specifically with respect to these peer-to-peer networks and applied the "inducement rule" to file-sharing services. Evidence demonstrated that 90 percent of the files available to download from Grokster and StreamCast were copyrighted works, and Grokster and StreamCast conceded that most users were downloading copyrighted material. There was also an abundance of evidence that through their respective software

applications and advertisements, both entities marketed themselves as the alternative to Napster, and their business models demonstrated "that their principal object[ive] was [the] use of their software to download copyrighted works."[126] The Court vacated the court of appeals' affirmation of summary judgment for Grokster and StreamCast, and rejected the court of appeals' broad interpretation of *Sony Corp. v. Universal City Studios*, but declined to further discuss the balance between protecting copyrighted works and promoting commerce in the context of how much noninfringing use each service was capable of providing, and did not at all discuss the issue of fair use. Instead, the Court noted that *Sony* did not preclude other forms of infringement liability and, focusing on the intent of the defendants in their inducement of file sharing, held that "one who distributes a device with the object of promoting its use to infringe copyright, as shown by clear expression or other affirmative steps taken to foster infringement, is liable for the resulting acts of infringement by third parties."[127] Citing *Sony*, the Court further opined that mere knowledge of potential or actual infringement are not sufficient bases for liability, but that "the inducement rule … premises liability on purposeful, culpable expression and conduct, and thus does nothing to compromise legitimate commerce or discourage innovation having a lawful purpose."[128]

Since the service and software in *Grokster* had other lawful purposes, the Supreme Court's decision underscores the importance of proving an intent to infringe or cause infringement. Thus, when asking a court to look behind stratagems and disclaimers that hide unlawful purposes, the copyright holder should consider what other evidence exists or is likely to exist of product design, advertising, marketing, external and internal communications, revenue plans, and other factors that would prove unlawful intent. In addition, for copyright holders, the problem remains that many providers of file-sharing software may not be subject to the jurisdiction of U.S. courts and that file-sharing software, such as "Darknet," provides anonymity to users illegally downloading copyrighted materials. As will be discussed, many countries are signatories to TRIPS (see Section 11.11.1 of this chapter) and subscribe to international copyright protection. Following Grokster, the maker of KaZaa file-sharing software was enjoined in Australia from using its software to commit copyright infringement. The remedy required alteration of the software so that it would not duplicate copyrighted works.

11.9 OTHER TOOLS TO PREVENT UNAUTHORIZED INTRUSIONS. Several legal principles and laws support the right to prevent and prosecute unauthorized intrusions. These include the definition of trespass, terms of use, and several critically important and widely used laws explicitly addressing the issues.

11.9.1 Trespass. Trespass is a common law concept that we are all familiar with when applied to land. We have all seen and probably at some point in our youth violated the no-trespassing signs that are posted on an unfriendly neighbor's property. Trespass is also a concept that can apply to computers and informational databases. Courts have been taking older concepts and reapplying them to new situations.

In *eBay, Inc. v. Bidder's Edge, Inc.*,[129] the Federal District Court granted eBay an injunction forbidding Bidder's Edge from using a software robot to scrape information from eBay's Web site. The court based the injunction on its finding that accessing the Web site in a manner that was beyond eBay's posted notice (there were actual letters of objection) constituted a trespass. The court reasoned that the "electronic signals

sent by Bidder's Edge to retrieve information from eBay's computer system [were] sufficiently tangible to support a trespass cause of action." The court further viewed the ongoing violation of eBay's fundamental right to exclude others from its computer system as creating sufficient irreparable harm to warrant an injunction. Thus, it was not necessary that eBay prove that the access actually interfered with the operation of the Web site. Rather, proof of the "intermeddling with or use of another's personal property" was sufficient to establish the cause of action for trespass. What is significant here is that eBay did permit others to access its Web site under license, and the court viewed conduct that exceeded the licensed use, upon notice to the violator, to be a trespass.

However, the applicability of trespass to unauthorized computer activity is not settled. Where trespass involves an object, rather than land, there must not only be improper use but also some harm to the physical condition or value of the object, or the misuse must deprive the rightful owner of the use of the object for a substantial period of time. The two must be causally related. In *Intel v. Hamidi*,[130] the California Supreme Court reversed a lower court's banning a former employee from sending unsolicited emails on the grounds of trespass. The court thought that the reach of the doctrine had been extended too far, concluding that bad analogies (i.e., viewing servers as houses and electronic waves as intrusions) create bad law. The court declined to view computers as real property. Rather, finding that they were like other personal property, the court found that this communication was no different from a letter delivered by mail or a telephone call. In short, the court declined to find a trespass because there was an "unwelcome communication, electronic or otherwise" that had fictitiously caused an "injury to a communication system." Here there was no injury to the computer system although Intel claimed injury to its business.

Intel v. Hamidi simply warns against overbreadth of application of the concept of trespass. If injury to the computer system can be demonstrated, then the concept of trespass does lie as a tool in the arsenal of remedies assuming that the trespasser can be identified.

11.9.2 Terms of Use. Terms of use can constitute a contract with respect to Web site usage. Thus, in any situation where electronic access is requested or permitted, the terms and conditions of use, together with an acknowledgment that such terms have been seen and consented to, can be enforced as restricting usage. In *Register.com, Inc. v. Verio, Inc.*,[131] the Second Circuit upheld an order enjoining Web site access primarily on the issue of contract. There, as described by the Second Circuit, the defendant Verio, against whom the preliminary injunction was issued, was engaged in the business of selling a variety of Web site design, development, and operation services. In the sale of such services, Verio competed with Register's Web site development business. To facilitate its pursuit of customers, Verio undertook to obtain daily updates of the WHOIS information relating to newly registered domain names. To achieve this, Verio devised an automated software program, or robot, which each day would submit multiple successive WHOIS queries through the port 43 accesses of various registrars. Upon acquiring the WHOIS information of new registrants, Verio would send them marketing solicitations by email, telemarketing, and direct mail. To the extent that Verio's solicitations were sent by email, the practice was inconsistent with the terms of the restrictive legend Register attached to its responses to Verio's queries.

Register at first complained to Verio about this use and then adopted a new restrictive legend on its Web site that undertook to bar mass solicitation "via direct mail, electronic

mail, or by telephone." The court concluded that Verio's conduct formed a contract, like buying an apple at a roadside fruit stand, which Verio breached:

> We recognize that contract offers on the Internet often require the offeree to click on an "I agree" icon. And no doubt, in many circumstances, such a statement of agreement by the offeree is essential to the formation of a contract. But not in all circumstances. While new commerce on the Internet has exposed courts to many new situations, it has not fundamentally changed the principles of contract. It is standard contract doctrine that when a benefit is offered subject to stated conditions, and the offeree makes a decision to take the benefit with knowledge of the terms of the offer, the taking constitutes an acceptance of the terms, which accordingly become binding on the offeree. See, e.g., *Restatement (Second) of Contracts § 69* (1)(a) (1981) ("Silence and inaction operate as an acceptance … where an offeree takes the benefit of offered services with reasonable opportunity to reject them and reason to know that they were offered with the expectation of compensation.")
>
> * * *
>
> Returning to the apple stand, the visitor, who sees apples offered for 50 cents apiece and takes an apple, owes 50 cents, regardless whether he did or did not say, "I agree." The choice offered in such circumstances is to take the apple on the known terms of the offer or not to take the apple. As we see it, the defendant in Ticketmaster and Verio in this case had a similar choice. Each was offered access to information subject to terms of which they were well aware. Their choice was either to accept the offer of contract, taking the information subject to the terms of the offer, or, if the terms were not acceptable, to decline to take the benefits.

Id., at 403; and was also a trespass because:

> The district court found that Verio's use **[**31]** of search robots, consisting of software programs performing multiple automated successive queries, consumed a significant portion of the capacity of Register's computer systems. While Verio's robots alone would not incapacitate Register's systems, the court found that if Verio were permitted to continue to access Register's computers through such robots, it was "highly probable" that other Internet service providers would devise similar programs to access Register's data, and that the system would be overtaxed and would crash. We cannot say these findings were unreasonable.

Id., at 405.

Similarly, although in a different setting, in *ProCD v. Zeidenberg*,[132] where ProCD sold a CD with noncopyrightable data. Access to the data, however, was controlled by a license agreement; if there was no acceptance, there was also no access. The license agreement prohibited the use of the data for any commercial use. Zeidenberg took the data and posted it on a Web site, which he used commercially to sell advertising. Thus, the data were being used to attract visitors. The court found the license limitation on use enforceable.

The importance of this decision is that so long as the owner prominently specifies the limitations, the restrictions can become a contract that is accepted by accepting the benefits of access and can be one safeguard against misuse of the access.

11.9.3 Computer Fraud and Abuse Act[133]

11.9.3.1 Prohibited Behavior and Damages. In 1984, Congress passed the original version of the Computer Fraud and Abuse Act (CFAA).[134] The general purpose was to protect "Federal interest computers" by criminalizing intentional and unauthorized access to those computers that resulted in damage to the computers or the data stored on them. The statute was substantially amended in 1986[135] and again in 1996[136] and now contains both criminal and private civil enforcement provisions.

The statute proscribes these activities:

… knowingly accessing a computer without authority or in excess of authority, thereafter obtaining U.S. government data to which access is restricted and delivering, or attempting to deliver, the data to someone not entitled to receive it[137];

intentionally accessing a computer without authority or in excess of authority and thereby obtaining protected consumer financial data[138];

intentional and unauthorized access of a U.S. government computer that affects the use of the computer by or for the U.S. government[139];

accessing a computer used in interstate commerce knowingly and with the intent to defraud and, as a result of the access, fraudulently obtaining something valued in excess of $5,000[140];

causing damage to computers used in interstate commerce by (i) knowingly transmitting a program, code, etc. that intentionally causes such damage, or (ii) intentionally accessing the computer without authority and causing such damage[141];

knowingly, and with the intent to defraud, trafficking in computer passwords for computers used in interstate commerce or by the U.S. government[142]; and

transmitting threats to cause damage to a protected computer with the intent to extort money or anything of value.[143]

The linchpin among the relevant decisions concerning access to data under the CFAA is whether the access is "without authority" or "in excess of authority." The factors considered by the courts include the steps taken by the owner of the information to protect against disclosure or use, the extent of the defendants' knowledge regarding their authority to access or use the data, and the use(s) made of the data after gaining access. The legislative history indicates that the statute was intended to "punish those who illegally use computers for commercial advantage."[144]

Broadly speaking, there are two sets of circumstances to consider. In the first instance, is the actual access authorized, either expressly or impliedly? In the Internet context, where there is a presumption of open access, the site or data owners must show that they took steps to protect the contents of their site and to limit access to the data at issue.[145] Once those steps are taken, the protection constitutes a wall through which even automated search retrieval systems may not go without express permission.[146] Without the wall, there must be some evidence of an intent to access for an impermissible purpose, as when Intuit inserted cookies into the hard drives of home computers.[147]

Second, has the authorized access been improperly exceeded? Generally speaking, those who use their permitted access for an unauthorized purpose to the detriment of the site or data owner have violated the CFAA. Examples include employees who obtain trade secret information and transmit it via the employer's email system to a competitor for which the employee is about to begin work[148]; using an ISP subscription membership to gain access to and harvest email addresses of other subscribers in order to transmit unsolicited bulk emails[149]; and using access to an employer's email system to alter and delete company files.[150]

The criminal penalties range from fines to imprisonment for up to 20 years for multiple offenses. As discussed in Section 11.9, the CFAA has become a prominent element of claims by the U.S. government and private parties seeking to protect data that are not always protected by other statutory schemes.

11.9.3.2 Its Application to Web Crawling and Bots. Web robots, or "bots," have become widespread to scrape data from Web sites. All of that data generally are available to the public. That is, any individual can access the same information,

but not with the speed or accuracy of a Webspider. But when does such "scraping" run afoul of the CFAA? To what extent does the law protect site operators or company data from penetration by an outside third party?

The key to the analysis under the CFAA is to ask whether the data are in fact publicly available. Are there technical barriers, such as passwords or codes, that have to be circumvented? Do the terms of use prohibit access or use other than by an individual consumer? These questions are critical to determining whether the access either exceeds authority or is without authority under the CFAA.

If the answer to either one of these questions (or similar questions) is yes, one needs to consider access carefully since such access and downloading of data is likely to violate the CFAA. In *EF Cultural Travel v. Zefer Corporation*, Zefer designed a Web bot to scrape travel trip and pricing information from the Web site of EF Cultural Travel (EF) for use by a competitive travel Web site. The bot, designed by Zefer, downloaded the information by calling URLs on which each separate trip and pricing information was stored, reading the source code for the key features, and storing the information on a spreadsheet. The bot did so in a fashion not to burden or interfere with EF's Web site. Once gathered, the information was turned over to a competitor, who used the information to adjust price and trip information that it offered. Zefer's scraping did not occur continuously, but only on two dedicated occasions. EF sued, claiming that a violation of the CFAA had occurred. The First Circuit Court of Appeals disagreed, refusing to read into what is or is not authorized some "reasonable expectations" standard, instead requiring that the Web site operator expressly state any limitations on access in its terms and conditions. On remand to the Federal District Court, the court, following the First Circuit, granted summary judgment for Zefer.

11.9.3.3 Simple Preventive Measures. Not surprisingly, there are several methods for preventing unauthorized access in the first instance and, if unsuccessful, in prevailing in any subsequent claim arising under the CFAA. Perhaps the most obvious measure, and one that the First Circuit Court of Appeals underscored, is to make sure that each visitor to a Web site is adequately notified that the owner of the site intends only limited use or access to the data on the site. The notice can take many forms.

For example, a detectable message easily identifiable on a home page warning visitors that the posted information is available only for viewing and not for use in any manner adverse to the host's interests would be sufficient. Understandably, most Web hosts are reluctant to post such a blatant limitation—it is not necessarily "good for business." For those interested in an equally effective but less direct message, an increasingly common practice is to compel site visitors to register before gaining access to links and other pages available through the home page. The more difficult the registration process, the greater the host's apparent intent to restrict access to, and use of, the information that will be accessible after registration is completed.

Those hosts that require the payment of money, some kind of membership, or an access agreement before providing access establish what, for purposes of statutes like the CFAA that criminalize unauthorized access, will most often be seen as providing sufficient notice of the limits of authorized access. In the case of membership sites, the presumption is that each registrant is prequalified and therefore authorized to view and use the more restricted data, at least for purposes consistent with the terms of access. Enforceable click-wrap access agreements establish not only notice of access limitations; they also secure each visitor's agreement to use the Web site and the data therein within the stated limitations.

Securing Web-based data against unauthorized use or users is, in some ways, antithetical to the information-sharing intent and purpose of the Web. In this regard, however, the question arising when we post information on the Web differs little from the question posed over the centuries regarding the extent to which each of us wants our competitors or adversaries to use our proprietary work against our interests. The greater the concern, the more likely that each host will have to limit the data posted on the Web, or else increase each visitor's awareness of the rules of access.

11.9.4 Electronic Communications and Privacy. Electronic privacy is becoming the issue in our society of databases and networking. Most of the U.S. "privacy" statutes are subject matter specific: the Telephone Consumer Protection Act of 1991 (do not call, for telemarketers); Health Insurance Portability Accountability Act of 1996 (privacy with respect to uses and disclosure of medical information); Children's Online Privacy Protection Act of 1998 (regulating collection of information from children under the age of 13 by Web sites directed to children); Gramm-Leach-Bliley Act of 1999 (regulating sharing of customer data by financial institutions); Controlling the Assault of Non-Solicited Pornography and Marketing Act of 2003 (restricting spammers and requiring an ability to opt out); the Fair and Accurate Credit Transaction Act of 2003 (providing very limited assistance with respect to identity theft such as the obligation to provide a yearly credit report). These laws do not provide assurances of privacy in the same way that the European Union did in its 1996 Data Protection Directive.[151] The EU Data Directive establishes protections against release of personal data, including emails, within the European Union, and restrictions on the transmission of such data outside the EU to countries or companies that do not have equivalent protections in place.[152]

In 2005, ChoicePoint, a large data broker, admitted that it had sold personal data on over 160,000 people to phony companies established by identity thieves. Since then, other companies have announced data break-ins and data leaks. As a result of such data security breaches, approximately half of the states have passed laws that require disclosure of unauthorized access to personal data.[153]

In the United States, the primary protection for privacy remains a lawsuit for tortious invasion of one's privacy. Because those rights are defined state by state, a review is beyond the scope of this chapter. However, most states recognize some form of the tort of invasion of privacy, and the tort has been recognized in the *Restatement (Second) of Torts* § 652, which courts reference as an authoritative source of the law. In general, the Restatement makes actionable (a) intentional intrusion, that is highly offensive to a reasonable man, into the seclusion of another's private affairs, (b) the public disclosure of private facts if such disclosure is highly offensive to a reasonable person, and is not a legitimate public concern, and (c) the appropriation for his own use or benefit of the name or likeness of another.

This chapter has already discussed the fiduciary obligation owed by employees to their employers with respect to confidential information. The development of the tort of privacy suggests that companies owe a similar obligation to their employees. Although slightly different in scope, but foreshadowing the growing body of law in this area, in *Remsburg v. Docusearch, Inc.*,[154] the New Hampshire Supreme Court was faced with a database company that had supplied information to a client that included a woman's personal information. The client used it to confront her and kill her. The New Hampshire Supreme Court held that the company had to act with "reasonable care in

disclosing a third person's personal information to a client." This decision is as yet an unanswered invitation to other courts.

On the federal level, the CFAA, of course, does address "unauthorized" access to computerized information. In addition, Congress has enacted some statutory regulations that specifically address electronic communications and privacy.

11.9.4.1 Wiretap Act and Electronic Communications Privacy Act. The Omnibus Crime Control and Safe Streets Act of 1968, generally referred to as the Federal Wiretap Act,[155] established the general parameters for permitted interception of communications by law enforcement. As originally crafted, the Wiretap Act covered only "wire and oral communications." In 1986, Congress enacted the Electronic Communications Privacy Act (ECPA),[156] which amended the Wiretap Act and created the Stored Wire and Electronic Communications and Transactional Records Act (Stored Communications Act or SCA) to "update and clarify federal privacy protections and standards in light of changes in computers and telecommunication technologies."[157] The SCA makes it unlawful to knowingly access a prohibited electronic communications service facility without authority, or in excess of authority, and for such public service provider to disclose information contained in such facilities. The ECPA allows a private plaintiff to bring a claim for knowing or intentional violation of the statute to recover actual damages or the statutory minimum of $1,000.

The 1986 amendment extended the Wiretap Act's coverage to include "electronic communications," which is defined as "any transfer of signs, signals, writing, images, sounds, data, or intelligence of any nature transmitted in whole or in part by a wire, radio, electromagnetic, photo-electronic or photo-optical system."[158] "Intercept" is defined as "the aural or other acquisition of the contents of any wire, electronic, or oral communication through the use of any electronic, mechanical, or other device."[159] Consequently, the Wiretap Act now makes it an offense to "intentionally *intercept* … any wire, oral, or *electronic communication*."[160] Thus, the definitions in the act now cover Internet transmissions such as emails or file transfers.

There is an important exception to this prohibition. Under the "consent of a party" exception, it is permissible to intercept communications where "one of the parties to the communication has given prior consent to such interception."[161] The requisite consent may be express or implied from the surrounding circumstances.[162] Furthermore, an employer may obtain consent by informing the employee of the monitoring practices in an employment contract or in an employee handbook.[163]

Under the "provider exception," a provider of electronic communication services "whose facilities are used in the transmission of a wire or electronic communication, [may] intercept, disclose or use that communication in the normal course of his employment while engaged in any activity which is a necessary incident … to the protection of the rights or property of the provider of that service."[164] This exception may allow an employer to lawfully intercept communications to detect an employee's unauthorized disclosure of trade secrets to third parties.[165]

11.9.4.2 Contemporaneous Transmission Requirement. The Wiretap Act only prohibits *interceptions* of electronic communications,[166] a term that has been more narrowly defined by the courts than the definition in the act might suggest. The definition of interception provides that an individual "intercepts" a wire, oral, or electronic communication "merely by *acquiring* its contents, regardless of when or under what circumstances the acquisition occurs."[167] In the context of this section, a serious question arises about the legality of intercepting electronic communications as they

were being transmitted and once they were stored, either temporarily or permanently. Although Congress intended to liberalize one's ability to monitor "wire communications" while it sought to make the monitoring of "electronic communications" more difficult,[168] courts have held that Congress intended to make acquisitions of electronic communications unlawful under the Wiretap Act "*only if* they occur *contemporaneously with* their transmissions"[169] and before they actually cross the finish line and become stored.[170] This is, of course, an interesting fiction when applied to Internet transmissions, which consist of packages that are broken up and passed from router to router as well as from temporary storage to temporary storage. It is a far cry from the interception of a telephone call. It may simply be that in applying the language of the statute, the courts are faced with applying it to a technology that was not really in existence when the statute was amended in 1986.

In recent years, the courts have attempted to apply the contemporaneous transmission requirement to various situations. For example, cookies used to recover personal data from visitors to a Web site constitute an interception of a contemporaneous electronic communication and a violation of the Wiretap Act.[171] Noting that electronic communications are generally in transit and in storage simultaneously, the court reasoned that users communicated simultaneously with the pharmaceutical client's Web server and with the software company's Web server and, thus, the information was acquired contemporaneously with its transmission.[172]

Where electronic transmissions are found in RAM or on the hard drive, they are stored communications and can be retrieved because they are outside of the Wiretap Act.[173] Similarly, an email that is recovered after it has been sent and received does not satisfy the contemporaneous transmission requirement and therefore has not been intercepted under the Wiretap Act.[174] Perhaps in response to these and other decisions, in 2001 Congress amended the Wiretap Act to apply the contemporary transmission requirement to wire communications that could not be retrieved, thereby permitting the recovery of stored wire communications.[175]

11.9.4.3 Konop v. Hawaiian Airlines, Inc. The *Konop* decision appears to be the most oft-cited case on the issue of "interception" under the Wiretap Act. Konop, the plaintiff, was an airline pilot who created and maintained a Web site where he posted bulletins critical of his employer, Hawaiian Airlines, Inc., and the airline union. Konop controlled access to his Web site by requiring visitors to log in with a user name and password and by creating a list of authorized users.

An officer of Hawaiian Airlines asked one such authorized user for permission to use his name to access the Web site. The officer logged on several times, and another officer, using the same technique, also logged on to view the information posted on Konop's bulletin. Konop eventually filed suit against Hawaiian Airlines, alleging that it violated the Wiretap Act when its officer gained unauthorized access to Konop's Web site.

The court first reiterated that the act only prohibits *interceptions* of electronic communications.[176] "Interception," the court held, requires that the party acquire the information contemporaneous with its transmission, and not while it is in electronic storage. In this case, the court concluded that the employer did not violate the Wiretap Act because the officers accessed an electronic communication located on an idle Web site, which did not satisfy the contemporaneous transmission requirement.[177]

11.9.5 Stored Communications Act. Unlike the Wiretap Act, the Stored Communications Act (SCA),[178] as its name suggests, establishes the limitations of access to stored communications (i.e., communications accessed *after* their

transmission).[179] Specifically, the SCA makes it unlawful to "intentionally [access] without authorization a facility through which an electronic communication service is provided … and thereby [obtain], [alter], or [prevent], authorized access to a wire or electronic communication while it is in electronic storage."[180] The SCA defines "electronic storage" as "(A) any temporary, intermediate storage of a wire or electronic communication incidental to the electronic transmission thereof; and (B) any storage of such communication by an electronic communication service provider for purposes of backup protection of such communication."[181] The SCA exempts from liability conduct "authorized … by the person or entity providing a wire or electronic communications service"[182] or "by a user of that service with respect to a communication of or intended for that user."[183]

11.9.5.1 Electronic Storage: Backup Files. The essential element that separates the SCA from the Wiretap Act is that the accessed communications reside in electronic storage. Therefore, the first question is what constitutes electronic storage. In *Theofel v. Farey-Jones,*[184] the United States Court of Appeals for the Ninth Circuit attempted to answer this question.

In *Theofel,* overzealous lawyers for Farey-Jones secured, through a subpoena issued to an ISP, emails sent and received by their opponents in the lawsuit, a company called Integrated Capital Associates (ICA). The subpoena requested from the ISP virtually every email ever sent or received by ICA and its employees. In response, the ISP posted a smattering of the emails on a Web site accessible to Farey-Jones and its lawyers. When ICA learned of these activities, it sued Farey-Jones for, among other things, violation of the SCA.

According to the court in *Theofel*, Congress recognized that users of ISPs have a legitimate interest in protecting the confidentiality of communications in electronic storage at a communications facility. Moreover, this legitimate interest cannot be overcome by fraud or by someone who knowingly exploits a mistake that permits access to what is otherwise protected. The court found that the use of the subpoena to access ICA's emails when it was reasonably plain, at least to counsel, that the subpoena was invalid, negated any apparent authority that Farey-Jones and its lawyers may have had to view ICA's emails.

Farey-Jones claimed that the ICA emails were not in "electronic storage" and therefore no violation of the SCA occurred. The court disagreed. As stated earlier, electronic storage exists when messages are stored on a temporary, intermediate basis as part of the process of transmitting the message to the recipient, and when messages are stored as part of a backup process. In this instance, the court found that the emails, which had apparently been delivered to their recipients, were stored by the ISP as part of its backup process for retrieval after initial receipt. Access to those emails was therefore protected by the SCA, which Farey-Jones and its lawyers violated.

11.9.5.2 Electronic Storage: Temporarily Stored Communications. Recent cases interpreting the meaning of "temporary, intermediate storage … incidental to" transmission of the communication have adhered to the letter of the law more than its spirit. In two cases involving the installation of cookies that were subsequently accessed by software companies for commercial gain, the courts have held that cookies are permanently (or at least indefinitely) installed in the consumer's hard drive and therefore cannot be considered "temporary, intermediate storage."[185] The *Doubleclick* decision also emphasized that the "temporary, intermediate storage"

element of the SCA means what it says, that is, the prohibited conduct involves only the unauthorized access to communications while they are being temporarily stored by an intermediate and does not include access to stored messages after they have been received.[186] In the context of an employer's right to examine an employee's emails, the employee will have no claim that an employer has violated the SCA when the employer opens emails sent or received by the employee once the email has been either received or discarded.[187]

11.10 OPEN SOURCE. With the continued proliferation of the Internet and computer software, the licensing, distribution, and use of open source code has gained publicity and added importance in the practice of intellectual property and computer security. "Open source" describes the distribution of computer code that is available (i.e., open) to all others and therefore allows computer programmers to read, apply, and modify the code, and also redistribute any changes.[188] The open source movement began with Richard Stallman's development of Gnu's Not UNIX (GNU), a freeware form of UNIX that was meant to be free software (free as in the freedom to use, modify, and distribute the software).[189] GNU's development created the first open source license, the General Public License (GPL). Linux, an open source–based operating system and an alternative to Microsoft Windows, experienced tremendous growth through its use of the GPL.[190] The prevalence of open source issues is evidenced by the 1998 formation of the Open Source Initiative (OSI), which not only promotes open source development and encourages its use by business[191] but also offers links to and information about most of the available open source licenses.

11.10.1 Open Source Licenses. The author of an open source code holds a copyright that operates as other copyrights do, but the code is released under a certain license on a nonproprietary basis. There are various types of open source licenses. The first open source license was the GPL, as described. It offers the broadest application of free software. In contrast, other licenses do not seek to perpetuate the free nature of a particular program. According to the Open Source Initiative, there are nearly 60 open source licenses now available for authors of source code,[192] all of which assert certain requirements of the software user.

11.10.2 GPL. Licensing under the GPL is premised on Stallman's idea of "copyleft," which basically uses copyright as a tool to ensure the continued free distribution of source code.[193] In other words, the GPL affords application, modification, and distribution rights to the copyrighted source code only if the user agrees that the distribution terms remain the same. This creates an endless chain of GPLs attached to future distributions of either the original or derived versions, regardless of their form.[194] This endless chain often is referred to as the GPL's "viral effect," as GPL-protected code multiplies from any modifications of original GPL-protected code.[195] The GPL applies not just to an originally protected software program but also to what it broadly defines as the "Program":

> [A]ny such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language.[196]

Moreover, although the GPL also states that independent and separate sections of a derivative work are not subject to the GPL's terms when they are distributed as separate

works, the GPL does apply when the user distributes those same independent and separate "sections as part of a whole which is based on the Program. … "[197] The broad application given to the program under the GPL further enhances the viral effect of the license.

Other provisions of the GPL require users who distribute verbatim copies of the source code to publish copyright notices, disclaim warranties, and provide copies of the GPL. In addition, the modifier/user must attach to any modifications a notice that the software was changed, must distribute or license the software free of charge to third parties, and must provide appropriate copyright notices, warranty disclaimers, and GPL terms and conditions. In sum, the GPL's sweeping terms not only seek to achieve the free software goals of the FSF but also to impact whether authors chose the GPL, and whether businesses utilize software subject to the GPL.

11.10.3 Other Open Source Licenses. The Berkeley Software Distribution (BSD) License and the Massachusetts Institute of Technology (MIT) License are very similar in that they both require copyright notices, disclaimers of warranties, and liability limitations. The BSD further prohibits contributors or similar organizations from endorsing the program and also requires a copy of the BSD's terms to be distributed with the software.

11.10.4 Business Policies with Respect to Open Source Licenses. The issue of whether distribution of a proprietary work that incorporates a small portion of GPL-protected code subjects that proprietary work to the terms of the GPL has never been litigated.[198] This is one risk of using open source software. Another risk is that failure to comply with the GPL's terms could lead to litigation.[199] For instance, MySQL sought to enjoin Progress Software Corporation from distributing MySQL's Gemini program without a GPL-compliant agreement.[200] Because there was a factual dispute as to whether Gemini was a derivative work or an independent work under the GPL, and because Progress stipulated that it disclosed Gemini's source code and would withdraw the end user license for commercial users, the court did not grant the injunction as to the GPL.[201]

Given the expanding use of open source, businesses need to develop comprehensive policies addressing their use of open source to avoid liability and publicly releasing their own proprietary technology.[202] Concerns generally involve license requirements regarding the distribution of the software and its modifications,[203] since those activities usually require the company to release the source code for any distributed modification, and modifications often terminate vendors' support agreements.[204] In addition, distributing unmodified open source as part of a proprietary program may require the company to release its own proprietary open source code.[205] It is more likely, however, that the company would be enjoined from distributing the open source or would have to pay damages.[206] These considerations should be addressed not only through company policy but also by choosing the best source code to use in programming, given the company's internal and external needs and the specific licensing requirements of that source code.

11.11 APPLICATION INTERNATIONALLY. Because the laws of the United States are the laws of just one nation among many, the enforcement of U.S. law and the protection of intellectual property rights in large part depend on international treaties. To the extent that the infringing acts or acts of piracy may be deemed to occur in the United States, or the infringers can be found in the United States, then the United States

has sufficient jurisdiction over these acts to enforce its laws. In other words, such actors can be sued directly in the courts of the United States for violation of the laws of the United States.

Apart from direct enforcement, international protection is usually a vehicle of bilateral agreements between the United States and individual countries or a function of international protocols or treaties to which the United States is a signatory. Thus, for example, the Paris Convention for the Protection of Industrial Property[207] establishes a system for recognizing priority of invention, but only among member countries. In addition, there is the Patent Cooperation Treaty (PCT), a multilateral treaty with more than 50 signatories. The PCT permits the filing of an international application that simplifies the filing process when a patent is sought in more than one nation. For copyright protection, there is also a series of international treaties and agreements that include the Berne Convention,[208] the Universal Copyright Convention, and the World Trade Organization (WTO) Agreement.[209] Canada, Mexico, and the United States also signed the North American Free Trade Agreement (NAFTA) in December 1992. NAFTA addresses intellectual property and requires that member states afford the same protections to intellectual property as members of the General Agreement on Tariffs and Trade (GATT). At a minimum, members of GATT must adopt four international conventions, including the Paris Convention and the Berne Convention.

These agreements, conventions, and treaties in large part do not attempt to reconcile the differences in the national laws of intellectual property. The particular national rules and nuances are simply too complicated, and there are too many differences of opinion to expect that these differences could be internally reconciled. Rather, in large measure, these international accords attempt to codify comity between the member nations so that each will recognize the legitimacy of the intellectual property rights in the other.

11.11.1 Agreement on Trade-Related Aspects of Intellectual Property Rights. On December 8, 1994, the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS) was signed into law in this country. The signing of TRIPS required changes to be made in United States statutes and regulations to bring them into conformity with international norms. TRIPS, however, was a product of the United States and other industrial countries pressing for stronger, more uniform standards for international treaties concerning intellectual property. The basic structure of TRIPS is to set the minimum standard of protection for intellectual property with each member nation free to adopt more stringent standards. Under the rubric used in the United States, TRIPS applies to copyrights, patents, trademarks, service marks, mask works (integrated circuit designs), and trade secrets. It also covers geographical indications[210] and industrial designs.[211] Not addressed by TRIPS, although part of the international jargon for intellectual property, are breeder's rights[212] and utility models.[213] Thus, TRIPS establishes no standards as applied to these concepts, leaving each nation to set the parameters of protection unimpeded by TRIPS.

It is not by accident that TRIPS was negotiated within the context of GATT, which had set the international standards for trade tariffs and had provided remedies of trade retaliation if such standards were not adhered to. The structure of GATT provided the means under which developing countries agreed to reduce their trade tariffs in exchange for the right to export innovative products under an exclusive monopoly conveyed by intellectual property rights. The second benefit to the GATT format was to provide a means for trade retaliation if, under the dispute resolution provisions of TRIPS, the WTO determines that there is noncompliance. In reality, it is obvious that TRIPS

benefits those industrial nations that are more likely to be at the forefront of innovation and more concerned with the protection of their citizens' intellectual property.[214] The major concession wrung by the developing countries under TRIPS was obtaining a period of 4 to 11 years to implement TRIPS and to bring their national laws into conformity.

TRIPS generally reflects the U.S. view that focuses on the economic underpinnings for intellectual property rights as serving the greater societal interests. There is thus a shift from "societal" interests to "enterprise" interests. In particular, TRIPS adopts high minimum standards for patents, which will require significant legislative changes in developing countries. The copyright section, however, affords less protection than may be afforded by European nations, but it is in line with treatment in the United States. In short, TRIPS responds to the concern of enterprises in the United States that too loose a system of international protection has enabled imitation of U.S. innovations through copying and outright piracy.

11.11.2 TRIPS and Trade Secrets. Under its category for "Protection of Undisclosed Information," TRIPS provides protection for the type of information routinely referred to as trade secrets in the United States. Member nations are required to implement laws that safeguard lawfully possessed information from being disclosed to, acquired by, or used by others without consent and contrary to "honest commercial practices" if such information is (a) a secret in that it is not in the public domain, (b) has commercial value because it is a secret, and (c) has been subject to reasonable steps to maintain its secrecy.

Because discussions that led to TRIPS are not institutionally preserved, unlike the United States Congressional Record, there is no negotiating history to be consulted to flesh out the meaning of the spare paragraphs instituting trade secret protection. There do, however, appear to be differences from the total panoply of protections afforded in the United States. The concept of public domain articulated by TRIPS is information that is "not, as a body or in the precise configuration and assembly of its components, generally known among or readily accessible to persons within the circles that normally deal with the kind of information in question." This articulation appears to be addressing technological formulations of information, as opposed to general commercial information, such as financial information, that is generally considered proprietary and confidential in the United States. The focus on a technology formulation for protected information is bolstered by the TRIPS requirement that the information have commercial value. Thus, other types of information that are not part of a traded article may be deemed to have no commercial value and therefore to fall outside of the scope of protection. Depending on the particular jurisdiction in the United States, there is a distinction between confidential information and trade secrets based on the requirement that a trade secret must have commercial value. This, in turn, has been held to mean that information that is not exploited commercially is unprotectable under the law of trade secret. For example, the results of failed experiments that never resulted in a commercial product lack commercial value, even though such experiments are certainly helpful in the next round of exploration, in that they are signposts of what not to do.

The lesson to be drawn is that one should not assume symmetry of protections just because of the TRIPS provision. Instead, as part of the reasonable steps to maintain secrecy, enterprises need to consider carefully thought out and structured contractual provisions as well as a system of data caching that leaves truly confidential data in the United States, even if access is permitted outside. Improper takings of such data are,

arguably, acts that occur in the United States, and such acts are subject to enforcement and punishment under the laws of the United States.

11.11.3 TRIPS and Copyright. TRIPS embraces the U.S. general model for copyright protection in its opening statement that "[c]opyright protection shall extend to expressions and not to ideas, procedures, methods of operation or mathematical concepts as such." All member nations agree that, as to the protection of copyrights, the Berne Convention will apply. Under the Berne Convention, the duration of a copyright is the life of the author plus 50 years. If the life of a natural person is not involved, then it is ordinarily 50 years from publication. In addition, computer programs, whether in source or object code, are to be protected as literary works under the Berne Convention. TRIPS also recognizes that compilations of data can be protected as creative works. Article 10, ¶ 2 explicitly provides:

> Compilations of data or other material, whether in machine readable or other form, which by reason of the *selection* or *arrangement* of their contents constitute *intellectual creations* shall be protected as such. Such protection, which shall not extend to the data or material itself, shall be without prejudice to any copyright subsisting in the data or material itself. (Emphasis added.)

TRIPS, therefore, does establish some minimum standard in the growing debate over what protections will be afforded a database. In the United States, the clear demarcation point for unprotected information is compilations that represent no more than "sweat-of-the-brow" efforts. Such compilations cannot be copyrighted.[215] The classic example of a sweat-of-the-brow effort is the copying and alphabetical organizing of names, addresses, and telephone numbers that are in telephone books. In the United States, the key for copyright protection is the creator's original contribution of selection and arrangement. Thus, arguably, the TRIPS provision mimics the law of the United States.

The European Union (EU) has taken a more protective path. In its 1996 European DataBase Directive, the EU granted databases *sui generis* protection as their own unique form of intellectual property. Under the EU Directive, a database is "a collection of independent works, data or other materials arranged in a systematic or methodical way and individually accessible by electronic or other means." A database may be protected either because it represents a work of "intellectual creation" or because it was compiled through "substantial investment." The EU Directive protects such databases from unauthorized extraction or use for a period of 15 years, with the ability to extend the period for an additional 15 years if there was a "substantial new investment" in the database. Such protection extends to databases of EU members and to databases of nationals of other countries that offer protections similar to the EU.

The United States, despite a number of legislative proposals, has not adopted a concomitant rule. The result, at least for multinationals, is that entities that rely on databases should consider "locating" such databases within an EU member to take advantage of the EU's database protections.

11.11.4 TRIPS and Patents. TRIPS requires that all members recognize the right to patent products or processes in all fields of technology. A patentable invention must be new, inventive, and have an industrial application. The patent application must fully and clearly disclose the invention so that a person skilled in the art could carry out the invention. The best mode for carrying out the invention as of the filing date must also be disclosed. Patent rights are to be enforced without discrimination as to

place of invention or whether the product is imported or produced locally. The patent of a product conveys the exclusive right to prevent, without consent of the inventor, the making, using, offering for sale, selling, or importing of the product. The patent of a process conveys the exclusive right to prevent all of the above for products that result from the process as well as the use of the process itself. The holder of a patent also has the rights to assign, transfer, or license the patent. The minimum period for a protecting a patent is 20 years from filing.

TRIPS gives each member state the right to carve out from patentability certain subject matters that have as their purpose the protection of human, animal, or plant life, or to avoid serious prejudice to the environment. In addition, TRIPS permits a member state to allow other use without authorization from the patent holder. The section defining when such use is permissible is the most detailed section among the patent provisions of TRIPS. In general, it permits such use only (a) after an effort to obtain a license from the patent holder on reasonable commercial terms and conditions, (b) with adequate remuneration to the patent holder, (c) if such use is limited predominantly to the domestic market of the member nation, and (d) if there is a review of the decision to permit, as well as the compensation, by a "higher authority in that Member."

One of the circumstances envisioned by TRIPS is the grant of a second patent that cannot be exploited without infringing an earlier (first) patent. In such cases, a member nation may grant authority if the invention embodied in the second patent represents an "important technical advance of considerable economic significance" with respect to the first patent's invention and a cross-license on reasonable terms is granted to the holder of the first patent to use the second patent. For process patents, TRIPS creates a limited burden on the alleged infringer to prove that the identical product was produced using a different process. In particular, a member state can create a presumption that the process patent was violated in circumstances where the product is new, or where the patent holder is unable to demonstrate what process was actually used.

11.11.5 TRIPS and Anticompetitive Restrictions. TRIPS acknowledges that some licensing practices or other conditions with respect to intellectual property rights may restrain competition, adversely affect trade, and impede the transfer and dissemination of technology. Accordingly, TRIPS permits member nations to specify practices that constitute an abuse of intellectual property rights and to adopt measures to control or limit such practices, so long as the regulation is consistent with other provisions of TRIPS. In the event that a national of a member nation violates another member's laws and regulations regarding anticompetitive activity, TRIPS provides for the right of the involved nations to exchange information confidentially regarding the nationals and their activities.

11.11.6 Remedies and Enforcement Mechanisms. Each member nation is expected to provide an enforcement mechanism under its national laws to permit effective action against any act of infringement. Such procedures are to include remedies to prevent acts of infringement as well as to deter future acts. TRIPS imposes the obligation that all such procedures be "fair and equitable" and not be "unnecessarily complicated or costly" or involve "unwarranted delays."[216] In general, these remedies mean access to civil judicial procedures with evidentiary standards that shift the burden of going forward to the claimed infringer, once the rights holder has presented reasonably available evidence to support its claim. Damages may be awarded sufficient to compensate the rights holder for the infringement if the "infringer knew or had reasonable grounds to know that he was engaging in infringing activity." This means

that vigilance and notice are essential to have meaningful protection for intellectual property rights, since notice is the best means for setting up a damage claim. TRIPS permits its members to allow the recovery of lost profits or predetermined (statutory) damages even when the infringer did not know that it was engaged in infringing behavior. Although injunctive relief is to be provided for, remedies may be limited in circumstances involving patent holders, as discussed, where adequate compensation is paid, and the alleged infringer has otherwise complied with the provisions of its national law permitting such use upon payment of reasonable compensation. In order to deter further infringement, infringing materials may be ordered destroyed or noncommercially disposed of.

In addition to civil remedies, TRIPS requires criminal penalties in cases of "willful trademark counterfeiting or copyright piracy on a commercial scale."[217]

11.12 RECENT DEVELOPMENTS IN INTELLECTUAL PROPERTY LAW[218]

11.12.1 AIA. Peter E. Heuser of Schwabe, Williamson, & Wyatt summarized the Leahy-Smith America Invents Act (AIA) of 2011[219] as follows: "The AIA is the most important legislative patent reform in over 50 years. The AIA will change how patents are granted, how patent litigation will proceed and what kinds of inventions are eligible for patents, among other things."[220] The author summarized the main features of the AIA in detailed discussions of the following areas:

- First-to-file will now establish priority of invention
- Prior commercial user defense is established
- New post-grant proceedings for patent validity challenges
- The Patent and Trademark Office (PTO) will no longer grant patents on tax strategy
- Special transitional review for certain patents related to financial products and services
- Most PTO fees will increase by 15 percent
- Limited prioritized examination will be available
- New rules will affect litigation by nonpracticing entities
- False patent marking claims are curbed
- Other provisions will make it more difficult to attack patent validity

Complete information about the legislation is available through the Library of Congress THOMAS database.[221]

11.12.2 The PROTECT IP Act (PIPA). The PROTECT IP Act (Preventing Real Online Threats to Economic Creativity and Theft of Intellectual Property Act)[222] or PIPA, was introduced in the U.S. Senate in May 2011 but failed to make it to the floor of the Senate.[223] After extensive public opposition, including a worldwide temporary blackout of thousands of Web sites in protest of PIPA and the Stop Online Piracy Act (SOPA, below),[224] the bill was suspended in January 2012 pending further analysis.[225]

PIPA's main points include the following (quoting several sections from the THOMAS database):

- Preventing Real Online Threats to Economic Creativity and Theft of Intellectual Property Act of 2011 or the PROTECT IP Act of 2011—(Sec. 3) Authorizes the Attorney General (AG) to commence: (1) an in personam action against a registrant of a nondomestic domain name (NDN) used by an Internet site dedicated to infringing activities (ISDIA) or an owner or operator of an ISDIA accessed through an NDN; or (2) if such individuals are unable to be found by the AG or have no address within a U.S. judicial district, an in rem action (against a domain name itself, in lieu of such individuals) against the NDN used by an ISDIA.
- Defines ISDIA as a site that: (1) has no significant use other than engaging in or facilitating copyright infringement, circumventing technology controlling access to copyrighted works, or selling or promoting counterfeit goods or services; or (2) is designed, operated, or marketed and used to engage in such activities.
- Defines NDN as a domain name for which the registry that issued the domain name and operates the relevant top level domain, and the registrar for the domain name, are located outside the United States.
- Allows the court, upon application by the AG after an NDN-related in personam or in rem action is commenced under this section, to issue a temporary restraining order or an injunction against the NDN, registrant, owner, or operator to cease and desist further ISDIA activity if the NDN is used within the United States to access an ISDIA directing business to U.S. residents and harming U.S. intellectual property right holders.
- Directs the AG to identify and provide advance notice to operators of nonauthoritative domain name system servers (NDNSSs), financial transaction providers (FTPs), Internet advertising services (IASs), and providers of information location tools (ILTs), including search engines, online directories, and other indexes with hypertext links or referrals to online locations, whose action may be required to prevent such NDN-related ISDIA activity.
- Sets forth the preventative measures required to be taken by NDNSSs, FTPs, IASs, and ILTs upon being served with a court order in such an NDN-related action commenced by the AG.
- (Sec. 4) Authorizes the AG or an intellectual property right owner harmed by an ISDIA to commence: (1) an in personam action against a registrant of an ISDIA's domain name or an owner or operator of an ISDIA accessed through a domain name; or (2) if such individuals are unable to be found or have no address within a U.S. judicial district, an in rem action against a domain name used by an ISDIA.
- Allows the court, upon application by the relevant plaintiff after an in personam or in rem action concerning a domain name is commenced under this section, to issue a temporary restraining order or injunction against a domain name, registrant, owner, or operator to cease and desist further ISDIA activity if the domain name is: (1) registered or assigned by a domain name registrar or registry located or doing business in the United States, or (2) used within the United States to access an ISDIA directing business to U.S. residents and harming U.S. intellectual property right holders.
- Directs the relevant plaintiff to identify and provide advance notice to FTPs and IASs whose action may be required to prevent such ISDIA activity.

- Requires, upon being served with a court order after such an in personam or in rem action concerning a domain name is commenced by the AG or a private right owner under this section: (1) FTPs to take reasonable specified preventative measures, and (2) IASs to take technically feasible and reasonable measures.
- Sets forth provisions regarding the entities that may be required to take certain preventative measures in actions concerning both domain names and NDNs: (1) granting immunity to such entities for actions complying with a court order, (2) authorizing the relevant plaintiff to bring an action for injunction relief against a served entity that knowingly and willfully fails to comply with a court order, and (3) permitting such entities to intervene in commenced actions and request modifications, suspensions, or terminations of related court orders.
- (Sec. 5) Provides immunity from liability for: (1) FTPs or IASs that, in good faith, voluntarily take certain preventative actions against ISDIAs, and (2) domain name registries and registrars, FTPs, ILTs, or IASs that, in good faith, withhold services from infringing sites that endanger public health by distributing prescription medication that is counterfeit, adulterated, misbranded, or without a valid prescription. …

11.12.3 The Stop Online Piracy Act (SOPA). The Stop Online Piracy Act (SOPA), H.R. 3261,[226] is summarized in the THOMAS database as follows:

- … Authorizes the Attorney General (AG) to seek a court order against a U.S.-directed foreign Internet site committing or facilitating online piracy to require the owner, operator, or domain name registrant, or the site or domain name itself if such persons are unable to be found, to cease and desist further activities constituting specified intellectual property offenses under the federal criminal code including criminal copyright infringement, unauthorized fixation and trafficking of sound recordings or videos of live musical performances, the recording of exhibited motion pictures, or trafficking in counterfeit labels, goods, or services.
- Sets forth an additional two-step process that allows an intellectual property right holder harmed by a U.S.-directed site dedicated to infringement, or a site promoted or used for infringement under certain circumstances, to first provide a written notification identifying the site to related payment network providers and Internet advertising services requiring such entities to forward the notification and suspend their services to such an identified site unless the site's owner, operator, or domain name registrant, upon receiving the forwarded notification, provides a counter notification explaining that it is not dedicated to engaging in specified violations. Authorizes the right holder to then commence an action for limited injunctive relief against the owner, operator, or domain name registrant, or against the site or domain name itself if such persons are unable to be found, if: (1) such a counter notification is provided (and, if it is a foreign site, includes consent to U.S. jurisdiction to adjudicate whether the site is dedicated to such violations), or (2) a payment network provider or Internet advertising service fails to suspend its services in the absence of such a counter notification.
- Requires online service providers, Internet search engines, payment network providers, and Internet advertising services, upon receiving a copy of a court order relating to an AG action, to carry out certain preventative measures including withholding services from an infringing site or preventing users located in

the United States from accessing the infringing site. Requires payment network providers and Internet advertising services, upon receiving a copy of such an order relating to a right holder's action, to carry out similar preventative measures.

- Provides immunity from liability for service providers, payment network providers, Internet advertising services, advertisers, Internet search engines, domain name registries, or domain name registrars that take actions required by this Act or otherwise voluntarily block access to or end financial affiliation with such sites.
- Permits such entities to stop or refuse services to certain sites that endanger public health by distributing prescription medication that is adulterated, misbranded, or without a valid prescription.
- Expands the offense of criminal copyright infringement to include public performances of: (1) copyrighted work by digital transmission, and (2) work intended for commercial dissemination by making it available on a computer network. Expands the criminal offenses of trafficking in inherently dangerous goods or services to include: (1) counterfeit drugs; and (2) goods or services falsely identified as meeting military standards or intended for use in a national security, law enforcement, or critical infrastructure application.
- Increases the penalties for: (1) specified trade secret offenses intended to benefit a foreign government, instrumentality, or agent; and (2) various other intellectual property offenses as amended by this Act.
- Directs the U.S. Sentencing Commission to review, and if appropriate, amend related Federal Sentencing Guidelines.
- Requires the Secretary of State and Secretary of Commerce to appoint at least one intellectual property attaché to be assigned to the U.S. embassy or diplomatic mission in a country in each geographic region covered by a Department of State regional bureau.

Critics of the legislation include the American Civil Liberties Association, some educators, some law professors, and the United States Student Association.[227] Arguments included the following:

- The bill would lead to removal of much noninfringing content from the Web, resulting in infringement of free speech.
- Eliminating the focus articulated in PIPA about concentrating on sites dedicated to infringing activity would waste government resources on an enormous range of sites.
- ISPs, search engine providers, payment network providers, and advertising services would all have to obey the Attorney General's orders to block all access to sites with infringing content, thus blocking access to all the sites' noninfringing content as well.
- Educational uses could be severely constrained if a single infringing document led to the shutdown of an entire site.
- Sites with a single link to infringing content could be classified as "facilitating" infringement and thus be shut down.

- The bill would violate standards of due process by allowing administrative shutdown without providing an opportunity for the owners of the accused sites a chance to defend themselves.
- SOPA's potential barriers to access could severely affect the worldwide movement to pressure dictatorial regimes such as that of the People's Republic of China in their consistent suppression of free access to information.
- Librarians, educators, and students could be subject to administrative shutdown even for what could be justified as fair use of copyright materials.

The proposed bill was dropped at the same time as PIPA (above).

11.12.4 Patent Trolls. Groups aggressively targeting users of little-known patents, often purchased from inventors who have never exercised their rights before, are known as *nonpracticing entities* or *patent trolls.* Some of these companies devote their entire business to suing or threatening to sue on the basis of their acquired patents.[228]

In one notorious case, a company bought

> … the Canadian patent known as "Automatic Information, Goods, and Services Dispensing System (Canada '216)" whose complete text is available at http://patents1.ic.gc.ca/ details?patent_number=1236216&language=EN_CA [and] specifically addresses "a system for automatically dispensing information, goods and services to a customer on a self-service basis including a central data processing centre in which information on services offered by various institutions in a particular industry is stored. One or more self-service information and sales terminals are remotely linked to the central data processing centre and are programmed to gather information from prospective customers on goods and services desired, to transmit to customers information on the desired goods or services from the central data processing centre, to take orders for goods or services from customers and transmit them for processing to the central data processing centre, to accept payment, and to deliver goods or services in the form of documents to the customer when orders are completed. The central data processing centre is also remotely linked to terminals of the various institutions serviced by the system, so that each institution can be kept up-dated on completed sales of services offered by that institution." [Note that Canadian spelling is used above.] Think about this patent. Does it not remind you unavoidably of what you did the last time you ordered a book or bought something online? Or performed any other commercial transaction on the Web?[229]

A study published by the Boston University School of Law[230] found that patent trolls "… cost U.S. software and hardware companies US$29 billion in 2011. …"[231]

In the House of Representatives, Peter DeFazio (D-OR) introduced HR.6245, Saving High-Tech Innovators from Egregious Legal Disputes Act of 2012, in August 2012.[232] It would "[Amend] federal patent law to allow a court, upon finding that a party does not have a reasonable likelihood of succeeding in an action disputing the validity or alleging infringement of a computer hardware or software patent, to award the recovery of full litigation costs to the prevailing party, including reasonable attorney's fees. …" At the time of writing (May 2013), the bill was still in the hands of the Subcommittee on Intellectual Property, Competition and the Internet of the House Committee on the Judiciary.

In May 2013, Senator Charles Schumer (D-NY) introduced S.866, the Patent Quality Improvement Act, an amendment to the AIA to extend its provisions for challenging patents on business methods.[233]

The Library of Congress THOMAS database describes the substance of the proposal as follows:

- Amends the Leahy-Smith America Invents Act to remove the eight-year sunset provision with respect to the transitional post-grant review program available to review the validity of covered business method patents, thereby making the program permanent.
- Expands the term "covered business method patent" to include a patent that claims a method or corresponding apparatus for performing data processing or other operations used in the practice, administration, or management of any enterprise, product, or service, except technological inventions. (Current law limits the program to financial products or services.)[234]

11.13 CONCLUDING REMARKS. Data security ultimately involves the protection of proprietary or personal data and intellectual property. The competition to acquire and retain intellectual property legally is invariably met by unethical and illegal efforts to deprive legitimate owners of their rights. It is necessary, therefore, to be fully aware of the mechanisms and procedures required to protect these rights as part of any computer security program.

This chapter has attempted to delineate the most important aspects of the problem. However, many facets of the legal questions remain unanswered or have been answered generally rather than in the context of a particular problem. Prudent guardians of intellectual property should monitor relevant judicial determinations continuously and be certain to integrate them into a planned approach to protect these most valuable assets.

11.14 FURTHER READING

Bently, L., and B. Sherman. *Intellectual Property Law*, 3rd ed. Oxford, UK: Oxford University Press, 2008.

Bloomberg. "Intellectual Property News" Web site. http://topics.bloomberg.com/intellectual-property

Bouchoux, D. E. *Intellectual Property: The Law of Trademarks, Copyrights, Patents, and Trade Secrets*, 4th ed. Cengage Learning, 2012.

Guardian Newspaper. *Intellectual property* archive of more than 500 recent articles. www.guardian.co.uk/law/intellectual-property

McJohn, S. *Intellectual Property: Examples and Explanations*, 4th ed. New York, NY: Aspen Publishers, 2012.

Nard, C. A., D. W. Barnes, and M. J. Madison. *The Law of Intellectual Property*, 3rd ed. New York, NY: Aspen Publishers, 2011.

Poltorak, A. I., and P. J. Lerner. *Essentials of Intellectual Property*, 2nd ed. Hoboken, NJ: John Wiley & Sons, 2011.

Stim, R. *Patent, Copyright & Trademark: An Intellectual Property Desk Reference*, 12th ed. Berkeley, CA: Nolo Press, 2012.

11.15 NOTES

1. For the uninitiated, a tort is a civil wrong (i.e., an act or failure to act that violates common law rules of civil society, and is distinguished from criminal wrongdoing).
2. *See ProCD, Inc. v. Zeidenberg*, 86 F.3d 1447 (7th Cir. 1996) (product could be returned if shrink-wrap terms were unacceptable).

3. *See Information Handling Services, Inc. v. LRP Publications, Inc.,* 2000 U.S. Dist. LEXIS 14531 (E.D. Pa., Sept. 20, 2000) (limit on unauthorized copies); *Hughes v. America Online, Inc.,* 204 F. Supp. 2d 178 (D. Ma. 2002) (enforcing forum selection clause).
4. *See LLAN Systems, Inc. v. Netscout Service Level Corp.,* 183 F. Supp. 328 (D. Mass. 2002) (click-wrap software agreement enforceable under Uniform Commercial Code as acceptance of an offer).
5. *See Motise v. America Online, Inc.*, 346 F. Supp. 2d 563 (S.D. N.Y. 2004) (user who logged on through another's account is bound by the terms of use even though not read).
6. 356 F.3d 393 (2d Cir. 2004).
7. *See Kewanee Oil Co. v. Bicron Corp.,* 416 U.S. 470, 473, 94 S. Ct. 1879, 40 L. Ed. 2d 315 (1974).
8. It is easy to confuse the notion of common law trade secret law with protection of confidential information. There is a distinction, however. At its core, trade secret law requires commercial application and utility, which is not true of confidential information that is generally protected as a matter of contract. For example, a failed experiment has no commercial utility and is not generally considered a trade secret, although it easily could be deemed confidential information.
9. The need to protect the information from general dissemination is what, in part, has given rise to the practice of Non Disclosure Agreements.
10. UTSA, 14 U.L.A. § 2(a).
11. *See* Trade Secrets Act, 18 U.S.C. § 1905; *see also* J. Michael Chamblee, J. D., *Validity, Construction, and Application of Title I of Economic Espionage Act of 1996*, 177 A.L.R. Fed. 609, *2 (2003) (hereinafter "Chamblee at __"). Other federal statutes, such as the National Stolen Property Act, 18 U.S.C. § 2314, were likewise of marginal utility in combating the rising problem of economic espionage. *See* Chamblee at *2.
12. Craig L. Uhrich, *Article: The Economic Espionage Act—Reverse Engineering and the Intellectual Property Public Policy*, 7 Mich. Telecomm. Tech. L. Rev. 147148-49 (2000/2001) (hereinafter "Uhrich at __"). Uhrich observes that the FBI investigated over 200% more economic espionage cases in 1996 than it had in 1994. *See* Uhrich at 151.
13. 18 U.S.C. §§ 1831, 1832.
14. *Id.*
15. 18 U.S.C. §§ 1831, 1832.
16. 18 U.S.C. §§ 1832 and 3571.
17. 18 U.S.C. § 1839 (3).
18. *United States v. Lange*, 312 F.3d 263 (7th Cir. 2002) (emphasis added).
19. 18 U.S.C. § 1839.
20. The 1980 Computer Software Copyright Act carved out for owners of computer programs a right to adapt, and for that purpose to copy, the program so that it functions on the actual computer in which it is installed. *See* discussion under the subheading "Derivative Works."
21. *See*, e.g., *Computer Management Assistance Co. v. Robert F. DeCastro, Inc*. 220 F.3d 396 (5th Cir. 2000) and *Engineering Dynamics, Inc. v. Structural Software, Inc.,* 26 F.3d 1335 (5th Cir. 1994).

22. Ideas, if protectable at all, are protected by patent.
23. The Copyright Act, 17 U.S.C. § 109(b).
24. The Copyright Act itself in sections 108 through 121 provides detailed limitations on the copyright owner's exclusive rights. These limitations are simply a matter of statutory construction. In addition, courts developed the doctrine of fair use in an effort to balance the rights of copyright owner and the public interest. That doctrine is now codified as part of the copyright statute in 17 U.S.C. § 107.
25. *See* the House Report No. 94-1476, 94th Cong., 2d Sess. 62 (1976) on the 1976 Act.
26. The Copyright Act, 17 U.S.C. §102(b).
27. 797 F.2d 1222 (3rd Cir. 1986).
28. 982 F.2d 693 (2d Cir. 1992).
29. 977 F.2d 1510 (9th Cir. 1992), amended, *Sega Enterprises Ltd. v. Accolade, Inc*., 1993 U.S. App. Lexis 78.
30. 977 F.2d at 1527–1528.
31. 975 F.2d 832 (Fed. Cir. 1992), *petition for rehearing denied*, 1992 U.S. App. Lexis 30957 (1992).
32. 79 F.3d 1532 (11th Cir. 1996).
33. 350 F.3d 640, 645 (7th Cir. 2003).
34. *Evolution, Inc. v. Suntrust Bank,* 342 F. Supp. 2d 943, 956 (D. Kan. 2004).
35. Compare *Micro Star v. Formgen, Inc*., 154 F.3d 1107 (9th Cir. 1998) (infringement found because copyrighted images displayed) with *Lewis Galoob Toys, Inc. v. Nintendo of America, Inc.,* 964 F.2d 965 (9th Cir. 1992) (no infringement although product compatible with Nintendo product).
36. 17 U.S.C. § 107.
37. 510 U.S. 569 (1994).
38. *Id.* at 577.
39. *Id.* at 580.
40. *See* the House Report No. 94-1476, 94th Cong., 2d Sess. 62 (1976) on the 1976 Copyright Act.
41. 17 U.S.C. § 117.
42. *Aymes v. Bonelli*, 47 F.3d 23 (2d Cir. 1995).
43. 17 U.S.C. § 901(a).
44. *MGM Studios Inc. v. Grokster, Ltd*., 545 U.S. 913, 930 (2005).
45. *Playboy Enterprises v. Frena*, 839 F. Supp. 1552 (M.D. Fla. 1993); *see also Sega Enterprises v. MAPHIA*, 857 F. Supp. 679 (N.D. Cal. 1994), and 948 F. Supp. 923 (N.D. Cal. 1996) (providing site for and encouraging uploading of copyrighted games was copyright infringement).
46. *Religious Technology Center v. Netcom On-line Communication Services*, *Inc.*, 90 F. Supp. 1361 (N.D. Cal. 1995).
47. 17 U.S.C. § 502.
48. 17 U.S.C. § 503.
49. 17 U.S.C. § 504.
50. 17 U.S.C. § 505.
51. 17 U.S.C. § 506.

52. 17 U.S.C. § 504(a).
53. *See Harris Market Research v. Marshall Marketing and Communications, Inc.,* 948 F.2d 1518 (10th Cir. 1991).
54. *See Regents of the University of Minnesota v. Applied Innovations, Inc.,* 685 F. Supp. 698, *aff'd*, 876 F.2d 626 (8th Cir. 1987) 698.
55. *Id.*
56. *See Cream Records, Inc. v. Jos. Schlitz Brewing Co.,* 754 F.2d 826 (9th Cir. 1985).
57. *See Eales v. Environmental Lifestyles, Inc.,* 958 F.2d 876 (9th Cir. 1992), *cert. den.* 113 S. Ct. 605.
58. *See Softel, Inc. v. Dragon Medical and Scientific Communications Ltd.,* 891 F. Supp. 935 (S.D. N.Y. 1995). Interestingly, in this case, the court also held that any increase in the infringer's profit may be considered when calculating the profit that must be disgorged to the license holder.
59. *See Harper & Row Publishers, Inc. v. Nation Enterprises,* 471 U.S. 539, 105 S. Ct. 2218 (1985); *Data General Corp. v. Grumman Systems Support Corp.,* 36 F.3d 1147 (1st Cir. 1994).
60. 17 U.S.C. § 504(c)(1).
61. *Id.*
62. 17 U.S.C. § 504(c)(2).
63. The theoretical nature of the relationship between actual and statutory damages is illustrated dramatically when the copyright owner demonstrates that the infringement was willful. *See Peer International Corp. v. Luna Records, Inc.,* 887 F. Supp. 560 (S.D. N.Y. 1995), where the music publisher's president willfully infringed licensed and unlicensed works and was assessed $10,000 for the licensed works, $15,000 for the unlicensed works, and $25,000 that the president used in derivative format without permission even though actual damages were $4,107. Presumably, this resulted from the court's attempt to find a way to punish the infringer since the statute makes no provision for punitive damages.
64. *See Central Point Software, Inc. v. Nugent*, 903 F. Supp. 1057 (E.D. Tex. 1995).
65. *See Walt Disney Co. v. Powell,* 897 F.2d 565 (D.C.Cir. 1990).
66. 17 U.S.C. § 1202(b).
67. *Universal City Studios, Inc. v. Reimerdes*, 111 F. Supp. 2d 294 (S.D. N.Y. 2000).
68. 17 U.S.C. § 1201(a).
69. 17 U.S.C. § 1201(a)(2).
70. 17 U.S.C. § 1201(a)(3).
71. *Id.*
72. 2000 U.S. Dist. LEXIS 1889 (W.D. Wash. January 18, 2000).
73. *Id* . at 19–21.
74. *Universal City Studios, Inc. v. Reimerdes,* supra note 67.
75. *Universal City Studios v. Corley*, 273 F.3d 429 (2nd Cir. 2002).
76. *Id*. at 446–447.
77. *Id.* at 450–451.
78. 307 F. Supp. 2d 521 (S.D. N.Y. 2004),
79. 320 F.3d 1317 (Fed. Cir. 2003), *writ of certiorari denied,* 539 U.S. 928 (2003).
80. 847 F.2d 255 (5th Cir. 1988).

81. *See*, e.g., *Davidson & Assocs. v. Jung*, 422 F.3d 630, 639 (8th Cir. 2005).
82. 17 U.S.C. § 1201(f)(1).
83. 17 U.S.C. § 1201(f)(2).
84. 17 U.S.C. § 1201(f)(4).
85. 17 U.S.C. § 1201(f)(3).
86. 381 F.3d 1178 (Fed. Cir. 2004), *cert. denied*, 544 U.S. 923 (2005).
87. 387 F.3d 522 (6th Cir. 2004).
88. *Storage Tech. Corp. v. Custom Hardware Eng'g & Consulting, Inc*., 421 F.3d 1307 (Fed. Cir. 2005).
89. 17 U.S.C. § 1201(i)(1).
90. 17 U.S.C. § 1201(j)(3).
91. 17 U.S.C. § 1201(j)(1).
92. 35 USC § 113 requires the submission of a drawing "where necessary for the understanding of the subject matter to be patented."
93. 35 U.S.C. §§ 283 and 284.
94. 127 S. Ct. 1746, 1757 (2007).
95. *Webster's Seventh New Collegiate Dictionary* (1967 ed.), p. 644.
96. 2007 CSI/FBI Computer Crime and Security Survey (hereafter the CSI/FBI Survey), pp. 12–13. Although the percentage of organizations reporting Internet abuse is down substantially since this chapter was first published, it nonetheless remains a source of substantial concern. In the same study, 26 percent of respondents reported phishing where the respondent was fraudulently identified as the sender; 25 percent reported misuse of instant messaging and unauthorized access to information; and 17 percent reported theft of customer and/or employee data.
97. SIIA Anti-Piracy 2007 Year in Review (www.siia.org/piracy/yir_2007.pdf). According to the SIIA, the source of the financial loss described in the text is the research firm IDC.
98. *See* Lamb and Rosen, *Global Piracy and Financial Valuation of Intellectual Property,* pp. 11.1–11.3.
99. "The subject matter of copyright … includes compilations." 17 U.S.C. § 103.
100. *Feist Publications, Inc. v. Rural Telephone Service Co., Inc.* 499 U.S. §§ 340, 361 (1991).
101. *Id*. at 350–351. *See* 17 U.S.C. §§ 101–103.
102. *Id.* at 344, 348–349. *See Ticketmaster Corp. v. Tickets.com, Inc*., 2000 U.S. Dist. LEXIS 12987 (C.D. Cal. Aug. 10, 2000), *aff'd,* 2001 U.S. App. LEXIS 1454 (9th Cir. Jan. 22, 2001).
103. *Feist Pub., Inc. v. Rural Tel.,* supra note 105, at 352–354, where the court rejected the so-called sweat-of-the-brow doctrine.
104. *Matthew Bender & Co., Inc. v. West Publishing Co.,* 158 F.3d 674, 682 (2d Cir. 1998) ("[t]he creative spark is missing where: (i) industry conventions or other external factors so dictate the selection that any person composing a compilation of the type at issue would necessarily select the same categories of information, or (ii) the author made obvious, garden-variety, or routine selections."). *See also Silverstein v. Penguin Putnam, Inc.* 368 F.3d 77, 83 (2d Cir. 2004).

105. 121 S. Ct. 2381; 150 L. Ed. 2d 500; 2001 U.S. LEXIS 4667; 69 U.S.L.W. 4567 (2001). Note: The party appealing to the Supreme Court is named first.
106. The court found interesting the publishers' decision not to assert a claim of transformative fair use. *Id.* at 2390. *See* Section 12.1.2.3.3 (transformative use section), supra.
107. 464 U.S. 417 (1984).
108. Transformative fair use was recently applied to the use of Rio devices, which permit individual users to download purchased MP3 music files to a hard drive and then play them either on the PC or a CD. These devices were analogized to the Betamax time shifting discussed in *Sony* and were upheld primarily on that basis. *See Recording Industry Association of America v. Diamond Multimedia Systems, Inc*., 180 F.3d 1072 (9th Cir. 1999).
109. 17 U.S.C. § 512(c).
110. 17 U.S.C. § 512(k).
111. 17 U.S.C. § 512(i).
112. 17 U.S.C. § 512(c)(1). *See ALS Scan, Inc. v. RemarQ Communities, Inc.,* 239 F.3d 619 (4th Cir. 2001), where the court of appeals determined what notice was sufficient to remove the safe harbor protection. *See also In re Aimster Copyright Litig.,* 252 F. Supp. 2d 634 (N.D. Ill. 2002), *aff'd* 334 F.3d 643 (7th Cir. 2003), for general discussion of this safe harbor provision, where Aimster had actual knowledge of the infringement by its users and therefore could not avoid liability under the safe harbor.
113. 907 F. Supp. 1361 (N.D. Cal. 1995). The *Netcom* decision predated the DMCA and provided part of the rationale and reasoning used by Congress in drafting and passing Title II of the DMCA. *See* House Rep. 105-551(I), at 11.
114. The church raised a question of fact about the impact of the ISP's activity on the church's potential market by asserting that the posting of the church's materials on the bulletin board discouraged active participation by existing and potential congregants. Therefore, the court could not find for the ISP as a matter of law.
115. 373 F.3d 544, 555 (4th Cir. 2004). The court went on to state, however, that an ISP "can become liable indirectly upon a showing of additional involvement sufficient to establish a contributory or vicarious violation of the Act. In that case, the ISP could still look to the DMCA for a safe harbor if it fulfilled the conditions therein."
116. *See, e.g., Playboy Enterprises, Inc. v. Frena*, 839 F. Supp. 1552 (M.D. Fla. 1993). The *Frena* decision, insofar as it holds the bulletin board service provider liable for infringement, has been expressly overruled by Title II of the DMCA. *See* House Rep. 105-551(I), at 11.
117. *Los Angeles Times v. Free Republic*, 2000 U.S. Dist. LEXIS 5669 (C.D. Cal. April 5, 2000). In the *Free Republic* decision, the court recognized the public benefit of posting articles for commentary and criticism but found that the initial postings contained little or no commentary that might transform the article into a new original work. *See also Video Pipeline, Inc.*, v. *Buena Vista Home Entm't, Inc.*, 342 F.3d 191, 199 (3d Cir. 2003), rejecting the fair use defense for an online distributor that made its own movie clip previews and used them as movie trailers by copying short segments of plaintiff's movies in part because the online distributor benefited from the infringement.

118. *Storm Impact, Inc. v. Software of the Month Club,* 13 F. Supp. 2d 782 (N.D. Ill. 1998).
119. There are various names for the components of the software programs that actually travel through the Web looking for data, including bots, crawlers, spiders, scrapers, and automated data retrieval systems.
120. *Kelly v. Arriba Soft Corp.,* 336 F.3d 811 (9th Cir. 2003).
121. 416 F. Supp. 2d 828, 838–846 (C.D. Calif. 2006).
122. 508 F.3d 1146 (9th Cir. 2007).
123. 239 F.3d 1004 (9th Cir. 2001), *aff'd*, 284 F.3d 1091 (9th Cir. April 3, 2002).
124. *See also UMG Recordings, Inc. v. MP3.com, Inc*., 92 F. Supp. 2d 349 (S.D. N.Y. 2000), where the district court held that storing recordings from purchased CDs on MP3.com's servers for retransmission to other users was infringement and not transformative fair use.
125. *Metro-Goldwyn-Mayer Studios, Inc. v. Grokster, Ltd.*, 545 U.S. 913, 125 S. Ct. 125 (2005).
126. *Id.* at 926.
127. *Id.* at 936–937.
128. *Id.* at 937.
129. 100 F. Supp. 2d 1058 (N.D. CA 2000).
130. 30 Cal. 4th 1342; 71 P.3d 296; 1 Cal. Rptr. 3d 32 (2003).
131. 356 F.3d 393 (2d Cir. 2004).
132. 86 F.3d 1447 (7th Cir. 1996).
133. 18 U.S.C. § 1030.
134. Pub. L. 98-474, codified at 18 U.S.C. § 1030.
135. Pub. L. 99-474.
136. National Information Infrastructure Protection Act of 1996, Pub. L. 104–294.
137. 18 U.S.C. § 1030(a)(1).
138. 18 U.S.C. § 1030(a)(2).
139. 18 U.S.C. § 1030(a)(3).
140. 18 U.S.C. § 1030(a)(4).
141. 18 U.S.C. § 1030(a)(5). *See Hotmail Corporation v. Van$ Money Pie, Inc*., 1998 WL 388389, 47 U.S.P.Q.2d 1020 (N.D. Cal. 1998).
142. 18 U.S.C. § 1030(a)(6).
143. 18 U.S.C. § 1030(a)(7).
144. Senate Rep. 104-357, pp. 7–8.
145. *Register.com, Inc. v. Verio, Inc*., 126 F. Supp. 2d 238 (S.D. N.Y. 2000).
146. *Id.*
147. *In Re Intuit Privacy Litigation*, 138 F. Supp. 2d 1272 (2001). *But see U.S. v. Czubinski*, 106 F.3d 1069 (1st Cir. 1997), where the court of appeals found that an IRS employee who accessed private tax information in violation of IRS rules but did not disclose the accessed information could not be prosecuted under 18 U.S.C. §030(a)(4) because he lacked an intent to deprive the affected taxpayers of their right to privacy.
148. *Shurgard Storage Centers, Inc. v. Safeguard Self Storage, Inc.,* 119 F. Supp. 2d 1121 (W.D. Wash 2000).

149. *America Online, Inc. v. LCGM, Inc.*, 46 F. Supp. 2d 444 (E.D. Va. 1998).
150. *U.S. v. Middleton*, 231 F.3d 1207 (9th Cir. 2000).
151. Council Directive 95/46, 1995 O.J. (L.281) 31–50 (EC).
152. As a result the United States negotiated with the EU the Safe Harbor Arrangement, administered by the Federal Trade Commission, under which a U.S. company can opt in to compliance with the EU Data Directive.
153. For an updated list, go to www.pirg.org/consumer/credit/statelaws.htm.
154. 149 N.H. 148, 816 A.2d 1001 (2003).
155. 18 U.S.C. §§ 2511(1)(a) and 2502(a).
156. Pub. L. No. 99-508, 100 Stat. 1848 (codified throughout scattered sections of 18 U.S.C.).
157. S. Rep. No. 99-541, at 1 (1986), reprinted in 1986 U.S.C.C.A.N. 3555, 3555.
158. 18 U.S.C. § 2510(12).
159. Id. § 2510(4).
160. 18 U.S.C. § 2511(1)(a) (emphasis added); *Konop v. Hawaiian Airlines, Inc.*, 302 F.3d 868, 875 (9th Cir. 2002) (*Konop*) (noting the legislative history of the ECPA indicates that Congress wanted to protect electronic communications that are configured to be private, such as email and private electronic bulletin boards).
161. 18 U.S.C.A. § 511(2)(d). One should note, however, that as a result of the Patriot Act, an order from a U.S. or state attorney general is sufficient to permit the government to install a device to record electronic transmissions for up to 60 days where related to an ongoing criminal investigation. The FBI has in its arsenal a program known as Carnivore that essentially tracks a target's online activity. Recently, Freedom of Information inquiries by the Electronic Privacy Information Center (EPIC, www.epic.org) suggests that the FBI has discontinued use of Carnivore because ISPs, in light of the PATRIOT Act, may be providing information regarding a user's internet traffic directly to the government.
162. *Griggs-Ryan v. Smith*, 904 F.2d 112, 117 (1st Cir. 1990) (holding consent may be implied where the individual is on notice of monitoring of all telephone calls).
163. Federal law allows states to enact their own wiretapping statutes provided that the state statutes are at least as strict as the federal counterpart. Lynn Bernabei, Ethical and Legal Issues of Workplace Monitoring of Employee Communications, 2003 WL 22002093, *2 (April 2003) (hereinafter "Bernabei at __"). Bernabei notes that most states have adopted statutes that mirror the federal statutes and that at least 10 states, including Massachusetts, require the consent of both parties before the employer can record a conversation. *Id.*
164. 18 U.S.C. § 2511(2)(a)(i) (Supp. 2003).
165. *Briggs v. Am. Air Filter Co.*, 630 F.2d 414 (5th Cir. 1980) (holding employer could monitor employee's communication "when [the] employee's supervisor [had] particular suspicions about confidential information being disclosed to a business competitor, [had] warned the employee not to disclose such information, [had] reason to believe that the employee is continuing to disclose the information, and [knew] that a particular phone call is with an agent of the competitor.").
166. 18 U.S.C. § 2511(1)(a).
167. *Konop*, 302 F.3d at 876 (emphasis added).
168. *Id.*

169. E.g., *Wesley Coll. v. Pitts*, 974 F. Supp. 375, 386 (D. Del. 1997) (holding that the act criminalizes only the interception of electronic communications contemporaneously with their transmission, not once they have been stored); *Payne v. Norwest Corp*., 911 F. Supp. 1299, 1303 (D. Mont. 1995) (holding the appropriation of voicemail or similar stored electronic message does not constitute an "interception" under the act); *Steve Jackson Games, Inc. v. United States Secret Service*, 36 F.3d 457, 461–462 (5th Cir. 1994) (holding that the government's acquisition of email messages stored on an electronic bulletin board system, but not yet retrieved by the intended recipients, was not an "interception" under the Wiretap Act).
170. *See United States v. Councilman,* 418 F.3d 67, 69–70 (1st Cir. 2005) (*en banc*).
171. *In re Pharmatrak, Inc*., 329 F.3d 9, 21 (1st Cir. 2003).
172. *Id.*
173. *United States v. Councilman*, 245 F. Supp. 2d 319 (D. Mass. 2003) (Wiretap Act count dismissed against email service provider who was charged with attempting to use electronic communications passing through his service for commercial gain).
174. *Eagle Investment Systems, Corp. v. Tamm*, 146 F. Supp. 2d 105, 112–113 (D. Mass. 2001).
175. USA PATRIOT Act § 209, 115 Stat. at 283; *Konop*, 302 F.3d at 876–878 ("The purpose of the recent amendment was to reduce the protection of voice mail messages to the lower level of protection provided other electronically stored communications.")
176. 302 F.3d at 876.
177. *Id.* at 879.
178. 18 U.S.C. § 2701 et seq.
179. *Bernabei* at *2.
180. 18 U.S.C. §§ 2701(a)(1), 2707(a) (emphasis added).
181. *Id.* § 2510(17), incorporated by 18 U.S.C. § 2711(1).
182. 18 U.S.C. § 2701(c)(1).
183. 18 U.S.C. § 2701(c)(2).
184. *Theofel v. Farey-Jones*, 359 F.3d 1006 (9th Cir. 2004).
185. *In re DoubleClick, Inc. Privacy Litigation*, 154 F. Supp. 2d 497 (S.D.N.Y. 2001) (*Doubleclick*); *In re Toys R US, Inc. Privacy Litigation*, 2001 U.S. Dist. LEXIS 16947 (N.D. Ca. 2001).
186. 154 F. Supp. 2d at 511–512.
187. *Fraser v. Nationwide Mut. Ins. Co*., 2003 U.S. App. LEXIS 24856, *19 (3rd Cir. 2003).
188. Jeanie Duncan Fallon, *Open Source Licenses: Understanding the General Public License, Technology Licensing Primer*, p. 248 (2d ed. 2001).
189. Richard Stallman, The GNU Project, available at www.gnu.org/gnu/thegnuproject.html. Stallman also started the Free Software Foundation (FSF) in 1985.
190. John C. Yates and Paul H. Arne, *Open Source Software Licenses: Perspectives of the End User and the Software Developer*, 25th Annual Institute on Computer & Internet Law, vol. 2, p. 104 (2005). It is estimated that thousands of programmers have contributed to Linux.

191. It is considered less extreme than the FSF, which basically advocates for an end to proprietary rights as applied to software.
192. *See* www.opensource.org/licenses/
193. This makes sense especially considering the FSF's vision of free software and its insistence on setting forth those views in the preamble of the GPL.
194. Section 2 of the GPL states: "You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License."
195. Fallon, at 250.
196. GPL Version 2.
197. *Id*. at 2(c).
198. Lori E. Lesser, *Open Source Software: Risks, Benefits, & Practical Realities in the Corporate Environment, Open Source Software: Risks, Benefits, & Practical Realities in the Corporate Environment*, p. 41 (2004).
199. *See id.*
200. *Progress Software Corp. v. MySQL AB,* 195 F. Supp. 328, 329 (D. Mass. 2002).
201. *Id.* The court also noted that MySQL did not demonstrate the likelihood of irreparable harm during the pendency of the case.
202. *See* Stuart D. Levi and Andrew Woodard, "Open Source Software: How to Use It and Control It in the Corporate Environment," *The Computer Lawyer,* vol. 21 (Aug. 8, 2004). "[A] policy needs to balance the benefits and competitive advantages of open source with the risks of using source code developed by parties with whom the company may not have a formal relationship."
203. *See* Yates andArne, *supra* n. 195, p. 107.
204. *See* Levi and Woodard, *supra* n. 207.
205. Also consider the fact that discovery in the course of litigation would also involve releasing proprietary codes, as IBM was forced to do for some of its products involved in the *SCO* litigation. Although discovery is obviously a different publication from that required under the GPL, it is an important issue to consider.
206. *Id.*
207. The Paris Convention was initially concluded in 1883 and updated in 1967. It is administered by the World Intellectual Property Organization, an agency of the United Nations. The Paris Convention has provisions that apply to patents, trademarks, service marks, industrial designs (similar to design patents), and unfair competition. Approximately 174 nations are now signatories to the Paris Convention.
208. Until the adoption of TRIPS, the Berne Convention was the other major international agreement. Like the Paris Convention, it is administered by the World Intellectual Property Organization. The Berne Convention, first adopted in 1886, has undergone a series of revisions. The Convention includes "every production in the literary, scientific and artistic domain whatever may be the mode or form of its expression." Berne Convention, Art. 2, ¶ 1. Essentially, it assures that a work protected within a member state will also be protected outside of the member state without being subject to discriminating formalities. The number of signatories to the Berne Convention is presently 165 nations.

209. The WTO effectively began operating on July 1, 1995, as a result of the 1994 Uruguay Round Agreements. The WTO replaces GATT (General Agreement on Tariffs and Trade), which had been in operation since 1950. Congress ratified the Uruguay Round Agreements in December 1994. The WTO has approximately 132 member nations. In 1995, the WTO and the World Intellectual Property Organization (WIPO) signed a joint agreement that provides, among other things, for cooperation in providing legal technical assistance and technical cooperation related to the TRIPS Agreement for developing country members of either of the two organizations. The WIPO has approximately 171 members and is responsible for international cooperation in promoting intellectual property protection around the world. In particular, it looks after various international conventions, such as the Paris Convention and the Berne Convention.
210. Geographical indications are marks or other expressions that state the country, region, or place in which a product or service originates.
211. Industrial designs protect the aesthetic look of the product and are similar but not identical to the United States notion of trade dress. Products may be afforded protection based on novelty or originality of design, depending on national law.
212. Breeder's rights confer protection on new and different plant varieties.
213. Utility models protect the manner in which a product works or functions and as such are different from industrial design, which protects only the aesthetics of the product. Generally, utility models address mechanical functioning, which in the United States is not protectable unless patentable. Thus, the innovation in the United States must be significant to warrant protection.
214. Until 1989, the developing countries largely refused to negotiate standards. Threats by the United States of trade sanctions under the United States Trade Act played a significant role in altering the positions of economically weaker developing countries. In particular, China, India, Taiwan, and Thailand were all investigated.
215. *Feist Publications v. Rural Telephone System*, 499 U.S. 340 (1991).
216. TRIPS, Article 41.
217. TRIPS, Article 61.
218. This section was written by M. E. Kabay.
219. Leahy-Smith American Invents Act of 2011, H.R. 1249, www.govtrack.us/congress/bills/112/hr1249
220. Peter E. Heuser, "Recent Developments in IP Law of Interest to Business Attorneys," Web site of Schwabe, Williamson & Wyatt, February 26, 2013, www.schwabe.com/showarticle.aspx?Show=12770
221. Library of Congress, 2011.
222. Library of Congress, 2012.
223. Preventing Real Online Threats to Economic Creativity and Theft of Intellectual Property Act of 2011, S. 968, www.govtrack.us/congress/bills/112/s968/text
224. Jonathan Weisman, "After an Online Firestorm, Congress Shelves Antipiracy Bills," *New York Times*. January 20, 2012, www.nytimes.com/2012/01/21/technology/senate-postpones-piracy-vote.html?_r=0
225. Trevor Timm, "After Historic Protest, Members of Congress Abandon PIPA and SOPA in Droves," *Electronic Frontier Foundation | Deeplinks* (blog), January

19, 2012. https://www.eff.org/deeplinks/2012/01/after-historic-protest-members-congress-abandon-pipa-and-sopa-droves

226. Library of Congress, 2011.
227. Electronic Frontier Foundation, Collection of documents about SOPA. *Electronic Frontier Foundation* Web site. December 12, 2011, https://www.eff.org/search/site/sopa
228. Mark Hachman, "Inside The Mind of A Patent Troll: If It's Legal, It Must Be OK," *readwrite.com*, May 13, 2013, http://readwrite.com/2013/05/03/inside-the-mind-of-a-patent-troll-if-its-legal-it-must-be-ok
229. M. E. Kabay, "PanIP has rights: PanIP has patents on e-commerce-related systems," *NetworkWorld*, April 15, 2003, www.networkworld.com/ newsletters/sec/2003/0414sec1.html
230. James E. Bessen and Michael J. Meurer, "The Direct Costs from NPE Disputes: Abstract," *Social Science Research Network* Web site, June 28, 2012, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2091210##
231. Loek Essers, "'Patent Trolls' Cost Tech Companies $29 Billion Last Year, Study Says," *PCWorld*, June 27, 2012, www.pcworld.com/article/258395/patent_trolls_cost_tech_companies_29_billion_last_year_study_says.html
232. Library of Congress, 2012.
233. Grant Gross, "Senator introduces legislation targeting patent trolls," *PCWorld*, May 01, 2013, www.pcworld.com/article/2037005/senator-introduces-legislation-targeting-patent-trolls.html
234. Library of Congress, 2013.

INTRODUCTION TO PART II

THREATS AND VULNERABILITIES

What are the practical, technical problems faced by security practitioners? Readers are introduced to what is known about the psychological profiles of computer criminals and employees who commit insider crime. The focus is then widened to look at national security issues involving information assurance—critical infrastructure protection in particular. After a systematic review of how criminals penetrate security perimeters—essential for developing proper defensive mechanisms—readers can study a variety of programmatic attacks (widely used by criminals) and methods of deception, such as social engineering. The section ends with a review of widespread problems such as spam, phishing, Trojans, Web server security problems, and physical facility vulnerabilities (an important concern for security specialists, but one that is often overlooked by computer-oriented personnel).

The chapter titles and topics in Part II include:

12. **The Psychology of Computer Criminals.** Psychological insights into motivations and behavioral disorders of criminal hackers and virus writers
13. **The Insider Threat.** Identifying potential risks among employees and other authorized personnel
14. **Information Warfare.** Cyberconflict and protection of national infrastructures in the face of a rising tide of state-sponsored and non-state-actor industrial espionage and sabotage
15. **Penetrating Computer Systems and Networks.** Widely used penetration techniques for breaching security perimeters
16. **Malicious Code.** Dangerous computer programs, including viruses and worms, increasingly used to create botnets of infected computers for spreading spam and causing denial of service
17. **Mobile Code.** Analysis of applets, controls, scripts, and other small programs, including those written in ActiveX, Java, and Javascript
18. **Denial-of-Service Attacks.** Resource saturation and outright sabotage that brings down availability of systems and that can be used as threats in extortion rackets by organized crime

19. **Social-Engineering and Low-Tech Attacks.** Lying, cheating, impersonation, intimidation—and countermeasures to strengthen organizations against such attacks, which have increased drastically in recent years
20. **Spam, Phishing, and Trojans: Attacks Meant to Fool.** Fighting spam, phishing, and Trojans—trickery that puts uninformed victims at serious risk of fraud such as identity theft
21. **Web-Based Vulnerabilities.** Web servers, and how to strengthen their defenses
22. **Physical Threats to the Information Infrastructure.** Attacks against the information infrastructure, including buildings and network media

CHAPTER 12

THE PSYCHOLOGY OF COMPUTER CRIMINALS

Q. Campbell and David M. Kennedy

12.1 INTRODUCTION **12·1**

12.2 SELF-REPORTED MOTIVATIONS **12·3**

12.3 PSYCHOLOGICAL PERSPECTIVES ON COMPUTER CRIME **12·4**

12.4 SOCIAL DISTANCE, ANONYMITY, AGGRESSION, AND COMPUTER CRIME **12·4**

- 12.4.1 Social Presence and Computer Crime 12·6
- 12.4.2 Deindividuation and Computer Crime 12·6
- 12.4.3 Social Identity Theory and Computer Crime 12·7
- 12.4.4 Social Learning Theory of Computer Crime 12·8

12.5 INDIVIDUAL DIFFERENCES AND COMPUTER CRIMINALS **12·9**

- 12.5.1 Antisocial and Narcissistic Personalities 12·9
- 12.5.2 Five-Factor Model of Personality and Computer Criminals 12·10
- 12.5.3 Asperger Syndrome and Computer Criminals 12·11
- 12.5.4 Internet Abuse and Computer Crime 12·12

12.6 ETHICS AND COMPUTER CRIME **12·14**

12.7 CLASSIFICATIONS OF COMPUTER CRIMINALS **12·16**

- 12.7.1 Early Classification Theories of Computer Criminals 12·17
- 12.7.2 Rogers's New Taxonomy of Computer Criminals 12·19
- 12.7.3 Hacktivists and Cyberterrorists: Hacking for a Cause 12·20
- 12.7.4 Dangerous/Malicious Insiders (DI/MI) 12·21
- 12.7.5 Virus Creators 12·22

12.8 RECOMMENDATIONS **12·24**

12.9 FURTHER READING **12·26**

12.10 NOTES **12·26**

12.1 INTRODUCTION. Symantec's Internet Security Threat Report for 2011 indicated an 81 percent increase in network attacks compared to 2010; this, coupled with a reported 187 million identities that were exposed due to outsider attacks, suggests that the threat of computer crime is growing to unprecedented levels.[1] For the most part, the industry has relied upon legal and technological solutions to reduce the risks to information security. An alternate approach is to target the human element;

to try to understand the psychological motivations behind those who would exploit these technologies and to design security system accordingly.[2] The main drawback to this approach is that the information security field has traditionally relied on outdated stereotypes of computer criminals, which have lead to convoluted, overgeneralized, and inaccurate portrayals of these individuals. Contributing to this is the industry's over-simplification of computer crime and its reliance on a generic, all-encompassing view of the computer criminal. The computer underground is a vast and varied landscape that comprises many different subgroups, some of which are infantile and benign, and others that are criminal and destructive. The purpose of this chapter is to identify the various subgroups of computer criminals and to examine their differing motivations from a psychological perspective. Using theoretical perspectives from social, personality, and clinical psychology, we will review current research on the various subsets of computer criminals, ranging from script kiddies to malicious insiders, and provide recommendations for addressing the problem of computer crime at its source.

The National Institute of Justice defines a *computer criminal* as any individual who uses computer or network technology to plan or perpetrate a violation of the law.[3] Although the term *computer hacker* is often used interchangeably with computer criminal, they are not synonymous. The term *hacker* was originally used as an umbrella term to refer to a computer programmer who changes or alters code (i.e., hacks) in a unique or unorthodox fashion to solve a problem or to enhance its use. Such interventions may be legal or illegal depending on the circumstances, intent, outcome, or use of the hacked program.

Although a computer criminal, or *cracker*, sometimes also referred to as a *malicious or criminal hacker*, may fall under this broad definition, these individuals typically alter or exploit technology for destructive purposes or financial gain rather than for benign or creative functions. Common examples of computer crimes include Web page defacements, creation and distribution of viruses, unauthorized access of technology, theft of information, distributed denials of service (DDoS), and so on.

In recent years, computer security analysts have reported that many computer criminals are moving away from the hacking-for-fun-and-notoriety mindset to hacking for profit.[4] More recently, activist groups (*hacktivists*) have used the Internet as a way of spreading their messages of social and political discord by engaging in digital harassment of their targets. According to the 2012 "Verizon Data Breach Report," hacktivist groups represent a significant threat to network security, accounting for the majority of data thefts occurring in 2011.[5,6]

Computer crime is an obvious financial and societal problem that shows no signs of slowing. Researchers suggest that computer attacks will continue to grow in frequency and sophistication as technology continues to evolve. More specialized threats to social networks, peer-to-peer networks (P2P), handheld mobile devices, and nontraditional hardware systems (e.g., networked gaming consoles and point of sale devices), have been identified in the wild with increasing regularity.[7] Douglas Campbell, president of the Syneca Research Group Inc., states that "the dominant threat to the United States is not thermonuclear war, but the information war."[8,9]

One solution that has been offered as an effort to slow this disturbing trend is to examine the motivations of computer criminals from a psychological perspective. Computer-crime researchers suggest that understanding the psychological motivations behind cyber criminals would aid in both cybercrime prevention and protection.[10,11] Generating a psychological profile of the various subtypes of computer criminals would aid in creating preventive initiatives as well as more effective countermeasures in the fight against computer crime. Since computer crime is not solely a technological

issue but one involving human agents, psychological theories regarding anonymity, aggression, social learning, and individual difference factors may enable us to better understand the behaviors and motivations of the computer criminal.

Information security consultant Donn Parker asserts that the creation of an effective malicious hacker profile still remains an elusive goal in the information security field.[12] Therefore, the goal of this chapter is to survey past and current literature surrounding the psychological motivations of computer criminals. Theories from criminology, as well as social, personality, and clinical psychology, will be presented in an attempt to explain some of the possible motivations behind computer criminals. Based on these psychological research studies, we will conclude by offering recommendations for attenuating computer crime from the perspective of the perpetrator.

12.2 SELF-REPORTED MOTIVATIONS. Perhaps the simplest approach to understanding the mindset of computer criminals is having the perpetrators describe their motivations in their own words. Using various self-reporting measures, including surveys, open-ended questionnaires, and first-person interviews, researchers have consistently found a number of common accounts used by computer criminals to explain and justify their illicit and sometimes harmful behaviors.[13,14,15,16]

According to sociologist Paul Taylor, computer criminals report that they are motivated by an interacting mix of six primary categories: addiction, curiosity, boredom, power, recognition, and politics.[17] Using a phenomenological-interpretive interview approach that emphasizes the interviewee's perception of reality, sociologist Orly Turgeman-Goldschmidt similarly found that computer criminals reported curiosity, thrill seeking, the need for power, and the ideological opposition to information restrictions among the motivations for their behaviors.[18] Taylor suggests that the extensive use of computers by these criminals may result from a combination of both compulsive behaviors and intellectual curiosity. From an outsider's perspective, an advanced computer user's need to meet the swiftly changing demands of the computer industry may appear to be an indicator of computer abuse, when in actuality the constant use of technology is a consequence of the field. A relentless curiosity and desire for technological improvement is often used by computer criminals as a motivation for their behaviors.[19] Anecdotal evidence has also suggested that the frustrations that result from restrictive computing environments (e.g., network or Internet filters), coupled with a lack of sufficient intellectual stimulation, contribute to some computer criminals' unauthorized access attempts. Some reformed computer criminals have indicated that once they were provided with more liberal access to technology, they were able to focus their skills on practical and legal endeavors rather than illicit undertakings.[20]

In one of the most ambitious efforts to understand the mindset of criminal hackers, computer security consultant and reformed computer criminal Raoul Chiesa and colleagues created the Hacker's Profiling Project (HPP).[21] The aim of the HPP was to utilize criminal profiling techniques to develop a comprehensive profile of criminal hackers. Chiesa developed a questionnaire that was judiciously distributed to known criminal hackers and asked questions regarding personal demographics, technological skill, criminal history, and social relationships. The research revealed again that some of the main motivations of criminal hackers are curiosity, proving their self-worth to themselves and others, and feelings of adventure. These relatively benign motivations are also coupled with feelings of anger, frustration, and rebellion against authority. For many of these individuals, the Internet is viewed as the great equalizer. Because of the reduced social context cues that guide face-to-face interactions, Internet users are judged more on their technological skills rather than their social skills, gender, or

ethnicity. For a more comprehensive looked at the results of the HPP, see the recently published book *Profiling Hackers*.[22]

Contrary to their stereotypical portrayals in the news media and in fiction, computer criminals appear to have wide-ranging social networks that exist in both their online and offline environments.[23,24] Taylor indicates that both the need for power and recognition by their peers may be motivating factors for some cybervandals. Computer criminals report feelings of enjoyment and satisfaction when they prove themselves better than system administrators and their peers. Communications researchers Hyung-jin Woo, Yeora Kim, and Joseph Dominick report in their analysis of Webpage defacements that 37 percent of the prank-related defacements contained messages that bragged or taunted the system administrators. Twenty-four percent of these types of defacements contained statements aimed at obtaining peer recognition and 8 percent contained boastful and self-aggrandizing verbiage.[25]

12.3 PSYCHOLOGICAL PERSPECTIVES ON COMPUTER CRIME. Although self-reporting analyses can give us some insight into the motivations behind computer criminals, these types of descriptive methodologies typically yield incomplete and sometimes inaccurate results. Unless the causes for our behaviors are obvious, our explicit or consciously held explanations for our actions are often misguided. Research has found that our behaviors are frequently controlled by subtle situational variables and implicit attitudes of which we are not typically aware, and may be distinctly different from the conscious mechanisms we use to explain our actions.[26] Therefore, our conscious justifications for our actions may be inaccurate if we are unaware of more subtle cognitive processes. The next section examines more empirically based psychological theories of aggression and deviance to gain a further understanding of the factors that may be influencing the behaviors of computer criminals.

12.4 SOCIAL DISTANCE, ANONYMITY, AGGRESSION, AND COMPUTER CRIME. Many acts of computer crime can be categorized as demonstrations of aggressive behaviors. For example, cracking into a company's Web server and defacing a Web page, or launching a DDoS attack on an organization's computer network, thereby crippling its Internet connection, are common malicious and aggressive acts engaged in by computer criminals. Social psychological theories on hostility and violence suggest that people are more likely to commit acts of aggression when the perpetrator of these acts is anonymous and the threat of retaliation is low.[27] Since cybervandals frequently use nicknames (*nicks* or *handles*), stolen accounts, and spoofed Internet Protocol (IP) addresses when they engage in illegal activities, their behaviors may be more aggressive than when they are more easily identifiable. Computer criminals are overly confident that their crimes cannot and will not be traced back to their true identities. Computer criminals who deface Web pages are so confident that they are anonymous that they regularly *tag* the hacked Website by leaving their handles and the handles of their friends, and in some cases, even their Internet email addresses and Web page links.[28]

Due to the relative anonymity of the Internet and the technical abilities of cybercriminals, which enable them to further obfuscate their identities, the resulting emotional distance may be another factor that contributes to increased aggression online. For example, it is an extremely difficult and tedious task to identify computer criminals who launch DDoS attacks against computer networks. The attacker plants denial-of-service (DoS) programs into compromised shell accounts controlled by a master client. The master client will instruct every slave DoS program to *cooperatively* launch an

attack at the victim's host at a configurable time and date. Thus, the DoS attacks are not launched by the criminal's own computer; rather, the attacks come from innocent networks that have been compromised by the cracker. This additional layer makes it all the more difficult for information-security professionals to locate the attack's perpetrator. Much like the Website vandals, DDoS attackers are also confident that the attacks will not be traced back to their actual identities. Frequently, DDoS attackers will even brag on Internet Relay Chat (IRC) channels and Twitter about how many host nodes they have compromised and against which domain they are planning to launch new attacks.[29]

Situational influences on behaviors and attitudes work on the Internet much as they do in the real world. However, computer criminals who commit aggressive acts against their innocent victims do not see the immediate consequences of their actions. The computer screen and increased social distance that characterize interactions online can act as an electronic buffer between the attacker and victim. Like the participants in psychologist Stanley Milgram's famous obedience experiment, computer criminals are physically and emotionally removed from their victims while they are committing their harmful actions.[30] They do not witness firsthand the consequences of their computerized attacks. Automated cracking and DDoS scripts, coupled with the lack of social presence in computer-mediated interactions, may make it easier to attack an entity that is not only emotionally and physically distant, but also depersonalized (e.g., a system administrator working for a large corporation).

Consistent with social psychologist Albert Bandura's theory of moral disengagement, individuals who engage in unscrupulous behaviors will often alter their thinking in order to justify their negative actions.[31] According to Bandura, most individuals will not commit cruel or illicit behaviors without first engaging in a series of cognitive justification strategies that allow the person to view those actions as moral and just. Immoral behaviors can be justified by comparing them to more egregious acts, minimizing the consequences of the actions, displacing responsibility, and by blaming the victim themselves. Criminologist Marc Rogers posits that computer criminals may rely on a number of these disengagement strategies in an attempt to reduce the dissonance associated with their malicious activities.[32]

Studies conducted by sociologists Paul Taylor and Orly Turgeman-Goldschmidt suggest that many computer criminals are, in fact, engaging in forms of moral disengagement.[33,34] Their interviewees report that computer crime is driven by a search for answers and spurs the development of new technologies. They further indicate that their electronic intrusions cause no real monetary harm or damage to the victims, and that larger corporations that can afford any financial losses that are incurred from their digital transgressions. Web page crackers will often criticize and publicly taunt the system administrators for not properly securing their computers, suggesting that the victims deserved to be attacked.[35] Rogers suggests that this victim-blaming strategy is likely the most common form of moral disengagement that is employed by computer criminals.[36]

A study conducted by information technology researcher Randall Young and colleagues confirmed that computer criminals have a morally distorted view of their deviant activities, enabling them to socially justify their digital exploits.[37] Self-identified computer criminals attending a computer conference reported significantly higher levels of moral disengagement than a control group of university students. The self-identified criminals strongly felt that their digital intrusions were actually helpful to the companies that they invaded and that their friends and families would not think negatively of them if they were caught engaging in illegal computer hacking.

12.4.1 Social Presence and Computer Crime. Social psychologist Sara Kiesler and colleagues found that during face-to-face (FTF) interactions, conversants implicitly attend to social context cues (e.g., facial expressions and intonations) and use them to guide their social behaviors. Since these social context cues are absent or reduced during computer-mediated interactions, digital communication may be more deregulated than FTF discussions.[38,39] Kiesler and Lee Sproull suggest that the absence of social-context cues in computer-mediated communication hinders the perception of and adaptation to social roles, structures, and norms.[40] The reduction of social-context cues in computer-mediated communication can lead to deregulated behavior, decreased social inhibitions, and reduced concern with social evaluation. The most common variables examined in their experiments were hostile language in the form of "flaming" (aggressive, rude, and often ad hominem attacks) and post hoc perceptions of group members (i.e., opinions formed after interacting with members). One empirical study found that group members communicating via computer-mediated communication were more hostile toward one another, took longer to reach decisions, and rated group members less favorably than comparable face-to-face groups.[41] Another experiment reported that there were 102 instances of hostile communication during computer-mediated interactions, compared to only 12 instances of hostile commentary during comparable FTF discussions.[42]

Based on Kiesler's findings, computer criminals may be engaging in hostile behaviors partly due to this reduction of available context cues. Crackers who harass and victimize system administrators and Internet users may be engaging in these antisocial activities due to the reduced attention to and concern with social evaluations. There are numerous anecdotal accounts of computer criminals "taking over" IRC channels, harassing people online, deleting entire computer systems, and even taunting system administrators whose networks they have compromised.[43] Their criminal and aggressive behaviors may be partially attributed to the reduced social context cues in computer-mediated communication and the resulting changes in their psychological states (i.e., deindividuation) while online.

12.4.2 Deindividuation and Computer Crime. Disinhibited behaviors have also been closely linked to the psychological state of deindividuation. Deindividuation is described as a loss of self-awareness that results in irrational, aggressive, antinormative, and antisocial behavior.[44,45] The deindividuated state traditionally was used to describe the mentality of individuals who comprised large riotous and hostile crowds (e.g., European soccer riots, mob violence, etc.). Social psychologist Phillip Zimbardo suggested that a number of antecedent variables, often characteristic of large crowds, lead to the deindividuated state. The psychosocial factors associated with anonymity, arousal, sensory overload, loss of responsibility, and mind-altering substances may lead to a loss of self-awareness, lessening of internal restraints, and a lack of concern for social or self-evaluation.[46]

The factors associated with deindividuation also appear to be present during some online activities. For instance, Internet users are relatively anonymous and often use handles to further obscure their true identities. Many of the Websites, software programs, and multimedia files that typify the computing experience are sensory arousing and in some cases can be overstimulating. The Internet can be viewed as a large global crowd that individuals become submersed in when they go online. It is possible that the physical and psychological characteristics associated with the Internet that make it so appealing may also lead individuals to engage in antisocial and antinormative behaviors due to psychological feelings of immersion and deindividuation.[47]

Deindividuation is brought about by an individual's loss of self-awareness, and psychological immersion into a crowd due to the aforementioned antecedents.[48] The aggressive, hostile, and antinormative actions of computer criminals may be linked to the denindividuated state. Zimbardo found that when participants were deindividuated, operationalized by anonymity, darkness, and loud music, they would administer higher levels of electric shocks to subjects, and for longer lengths of time, than individuated participants. Like Zimbardo's participants, computer criminals may be engaging in hostile and aggressive behavior due to deindividuation—that is, as a direct result of anonymity, subjective feelings of immersion, and the arousing nature of computer and Internet use.[49,50]

12.4.3 Social Identity Theory and Computer Crime. Social psychologists Martin Lea, Tom Postmes, and Russell Spears have recently developed a social identity model of deindividuation effects (SIDE) to explain the influence of deindividuating variables on behaviors and attitudes during computer-mediated communications.[51,52,53] According to social identity theory, an individual's self-concept resides on a continuum with a stable personal identity at one end and a social identity at the other. Depending on whether the social self, usually in group situations, or individual self is salient, the beliefs, norms, and actions associated with that particular self-concept will have the greatest influence on the individual's actions and attitudes.[54] When one of our social identities is salient, the norms associated with that group identity tend to guide and direct our behaviors.

According to the SIDE model, the isolation and visual anonymity that characterizes our online environment serves to enhance our social identities. This increase in social identification with a group may polarize our behaviors and attitudes toward the prevailing norms of that collective.[55,56,57] Contrary to popular media stereotypes, computer criminals appear to have large social networks and frequently form groups and friendships with other like-minded individuals.[58] The use of handles and pseudonyms by these individuals combined with their physical isolation from each other may increase their aggressive and criminal tendencies depending on the overall norms associated with their online social groups. If the criminal collective values electronic intrusions and defacements more than programming and coding, then these behaviors will be exhibited to a greater extent by members who strongly identify with that group.

According to Henri Tajfel and John Turner's social identity theory (SIT), we tend to identify with ingroups, or those with whom we share common bonds and feelings of unity.[59] We have a bias toward our own group members and contrast them with outgroups, whom we perceive as different from those in our ingroup. While this ingroup bias or favoritism may benefit and protect our self-concepts, it may cause us to dislike and unfairly treat outgroup members (e.g., network administrators).

Communications researcher Hyung-jin Woo and colleagues used SIT to explain the motivations behind some Web page defacements.[60] SIT predicts that when groups are in competition for scarce resources, or feel threatened by outgroup members, there is a tendency for groups to respond aggressively toward each other. Ingroup members see improvements in collective self-esteem and enhanced feelings of group unity when they engage in attacks against outgroup members. Based on these predictions, Woo and colleagues hypothesized that computer criminals who are motivated by outgroup threats will express more aggressive and varied communication in Web page defacements than nonthreatened defacers.[61]

A content analysis of 462 defaced Web pages indicated that the majority of the defacements (71 percent) were classified as nonmalicious pranks. The most common

motivations behind these prankster attacks were to beat the system or its administrator, to gain peer recognition, to brag about accomplishments, and for romantic purposes. Twenty-three percent of the defacements were classified as militant attacks. The motivations behind these attacks were to promote groups associated with nationalism, ethnicity, religion, freedom of information, and anti-pornography.[62] Consistent with the predictions made by SIT, the militant attacks were characterized by significantly more varied content, obscene language, insults, severe threats, and violent images. SIT may be useful in predicting the frequency and severity of attacks by computer criminals. Although only a minority (23 percent) of Web page defacements in this study resulted from intergroup conflict, those attacks *were* more severe in nature.[63]

12.4.4 Social Learning Theory of Computer Crime. Criminologist Marc Rogers suggests that social learning theory (SLT) may offer some insight into the behavior of computer criminals.[64] According to psychologist Albert Bandura, individuals learn behaviors by observing the actions of others and their associated consequences.[65] SLT draws from B. F. Skinner's operant-conditioning model of learning where behaviors are learned or extinguished through schedules of reinforcement and punishment. However, Bandura's theory suggests that social learning occurs when an individual simply observes others' behaviors and reinforcements and forms a cognitive association between the two actions. Once the behavior is acquired, the learned actions are subject to external reinforcement, as in operant conditioning or in self-directed reinforcement.[66] According to the Social Structure and Social Learning model (SSSL), criminals learn deviant behaviors from their associations and subsequent imitations of deviant peers.[67] Through these peer associations, individuals learn to rationalize criminal behaviors. As Bandura suggests, if these criminal activities are rewarded (e.g., money, increased status, etc.), then the behavior will be strengthened.

Recently, there has been a growing amount of social and media attention focused on information security and computer criminals. Newspapers, magazines, and electronic news sources have reported thousands of incidents, interviews, and commentary related to computer crime. A number of these articles appear to glamorize hacking and the Internet underground.[68] The articles compare computer criminals to rock-and-roll superstars, James Bond-like spies, and international freedom fighters. Motion pictures and television shows like *The Matrix Trilogy, Mission Impossible, Hackers, Swordfish,* and *The X-Files* have all bestowed mythical qualities on rebellious computer criminals, while media outlets report computer criminals being recruited for high-paying government and industry jobs.[69] The media's glorification and glamorization of hacking and computer criminals, teaches some individuals that it pays to commit computer crime—at least, from a social learning perspective.

Many crimes involving computers are difficult to investigate and prosecute. The public learns via the media that computer criminals often are afforded fame and notoriety among their peers, and in the information security field, for their illegal activities. There are very few instances of computer criminals' being convicted and serving jail time as a consequence of their actions; usually the criminals are given light sentences. Their notoriety leads to media interviews, book and movie deals, even consulting and public speaking jobs.[70] Once an action is learned, SLT states that the behavior will be maintained via self-directed and external reinforcement. If computer criminals are rewarded for their illegal activities via the acquisition of knowledge and their elevated status in the hacker community, and the popular media continues to glamorize and focus on the positive consequences associated with computer crime, then the cost and prevalence of these illicit actions will continue to grow. Lending empirical support

to the social learning model, criminologist Thomas Holt found that the four tenets of the SSSL model reliably predicted cyberdeviance in a college population. Students who associated with cybercriminals, imitated cyberdeviants, held morally ambiguous definitions regarding cybercrime, and had their computer-based deviance reinforced, were more likely to engage in online criminal activities.[71]

SLT may offer one explanation for the illegal behaviors of computer criminals, especially the marked increase in recent years.[72] Instead of focusing on the supposedly positive consequences of computer crime, media outlets should stress the negative repercussions of computer crime for both the victims and the perpetrators. Social learning theorists would suggest modeling *appropriate* use of computers and the immediate negative ramifications of cyberdeviance as one element for fighting computer crime.

12.5 INDIVIDUAL DIFFERENCES AND COMPUTER CRIMINALS. Although situational factors can account for some of the behaviors of some computer criminals, one must not discount the impact of personality factors on their illicit activities. Attitudes and behaviors are often the product of both situational influences and individual personality traits.[73] It should be noted that there are few empirical studies that engage in a scientific examination of the personality traits of computer criminals, so without concrete evidence, the anecdotal claims regarding pathological traits of cybercriminals should be interpreted with caution. In addition, simply having traits that are *consistent* with a psychological disorder does not mean that one actually *has* the disorder.

12.5.1 Antisocial and Narcissistic Personalities. According to M. E. Kabay, some computer criminals exhibit insincerity and dishonesty in combination with superficial charm and an enhanced intellect, traits that are consistent with the *Diagnostic and Statistical Manual of Mental Disorders IV* (DSM-IV) criteria for *antisocial personality disorder*.[74] He also notes that some computer criminals commit their illegal behavior for little or no visible rewards despite the threat of severe punishment.

Another central characteristic of antisocial personality disorder is lack of clear insight by perpetrators regarding their behaviors.[75] Researchers have noted that computer criminals do not view their criminal actions as harmful or illegal.[76,77] These criminals sometimes rationalize or externalize their behaviors by blaming the network administrators and software designers for not properly securing their computers and programs.

Computer crime researchers Eric Shaw, Keven Ruby, and Jerrold Post also have suggested that some computer criminals demonstrate personality characteristics consistent with some elements of *narcissistic personality disorder*.[78,79] According to DSM-IV criteria, narcissistic individuals are attention seekers with an exaggerated sense of entitlement.[80] Entitlement is described as the belief that one is in some way privileged and owed special treatment or recognition.

Shaw and associates suggest that entitlement is characteristic of many "dangerous insiders," or information technology specialists who commit electronic crimes against their own organizations.[81] When corporate authority does not recognize the work or achievements of an employee to their satisfaction, the criminal insider seeks revenge via electronic criminal aggressions. Anecdotal evidence suggests that outside network intruders also may demonstrate an exaggerated sense of entitlement, as well as a lack of empathy for their victims, also characteristic of narcissistic personality disorder.

One self-identified computer criminal states, "we rise above the rest, and then pull everyone else up to the same new heights … We seek to innovate, to invent. We, quite seriously, seek to boldly go where no one has gone before."[82] Narcissistic individuals

also frequently engage in rationalization to justify and defend their behaviors.[83] Computer criminal Toxic Shock writes, "We are misunderstood by the majority. We are misunderstood, misinterpreted, misrepresented. All because we simply want to learn. We simply want to increase the flow of knowledge, so that everyone can learn and benefit."[84] Although it would be a mistake to generalize these hypotheses to the entire population without any empirical support, certain subsets of computer criminals may demonstrate characteristics that are consistent with aspects of both narcissistic and antisocial personality disorders.

12.5.2 Five-Factor Model of Personality and Computer Criminals. In one of the rare empirical studies looking at computer criminals, criminologist Marc Rogers examined the relationship between the five-factor model of personality and self-reported "criminal computer activity."[85] The five-factor model formulated by psychologists Robert McCrae and Paul Costa in 1990 suggests that an individual's personality can be accurately described using five core dimensions: extraversion (e.g., sociable), neuroticism (e.g., anxious), agreeableness (e.g., cooperative), conscientiousness (e.g., ethical), and openness to experience (e.g., nonconforming).[86]

Rogers hypothesized that individuals engaging in computer crime would demonstrate higher levels of:

- exploitation,
- hedonistic morality,
- manipulation,
- antagonism,
- undirected behaviors,
- introversion,
- openness to experiences, and
- neuroticism

than noncriminals would. Three hundred eighty-one psychology students from an introductory psychology class were administered the computer-crime index (CCI), which is a self-report measure of computer-crime activity, along with measures of exploitation, manipulation, moral decision making, and a five-factor personality inventory. Contrary to the researcher's expectations, individuals who committed computer crimes did not significantly differ from the nonoffenders on any of the five-factor personality measures. However, students who reported engaging in illegal computer activities did demonstrate more exploitive and manipulative tendencies.[87]

In contrast, in a follow-up study Rogers did find that extraversion was a reliable predictor of computer crime behavior. The less extraverted an individual (i.e., more introverted), the more likely they were to engage in illicit computing activities. He suggests that the differences in populations between the two samples, the latter being a Canadian liberal arts college and the former being U.S. students from a technology program, may have contributed to the discrepant findings.[88] It is clear that further research will need to be conducted in this area in order to clear up the discrepant findings. One possible limitation in both of these studies is that the questionnaires were administered using a pencil-and-paper format, which assesses attitudes and traits when the participants' *offline* identities are salient.[89] Internet researchers have long suggested that there is a distinct difference between our online and offline identities.[90] Therefore,

had the participants been administered measures in an electronic format, when their online identities were more salient, Rogers might have found differing results

12.5.3 Asperger Syndrome and Computer Criminals. Recently researchers have suggested a possible link between criminal hacking and a relatively new developmental disorder named *Asperger syndrome* (AS).[91,92,93] AS is a disorder that resides at the mild end of the pervasive developmental disorder (PDD) spectrum, with classic autism at the more severe end. PDDs are characterized by primary developmental impairments in language and communication, social relations and skills, as well as repetitive and intense interests or behaviors.[94] Unlike autism, individuals who are diagnosed with AS have higher cognitive abilities and IQ scores ranging from normal to superior. Individuals with AS also have normal language and verbal skills, although there are noticeable deficits in social communication. Those diagnosed with AS typically have severe and systematic social skill impairment or underdevelopment, difficulties with interpersonal communication, and repetitive patterns of interests, behaviors, and activities.[95]

According to clinical psychologist Kenneth Gergen, AS individuals must demonstrate social impairment. They may have a lack of desire or inability to interact with peers, and may engage in inappropriate or awkward social responses.[96] These individuals may have extremely limited or focused interests, and are prone to engage in repetitive routines. Although language development is often normal, these individuals may demonstrate unusual speech patterns (e.g., rate, volume, and intonation). Individuals with AS also may demonstrate clumsy motor behaviors and body language, as well as inappropriate facial expressions and gazing.[97]

One of the noted features of AS that is anecdotally linked to computer criminals is the obsessive or extremely focused area of intellectual interest that the individuals demonstrate. Children with AS often show a preoccupation in areas such as math, science, technology, and machinery. They strive to learn and assimilate as much information as possible about their specialized interest. Researchers have indicated that their preoccupation may last well into adulthood, leading to careers associated with their intellectual interests.[98] Much of their social communication is egocentric, revolving around their obsessive interests, often leading to strained and difficult social interactions. Although children with AS desire normal peer interaction, their egocentric preoccupations, lack of appropriate social behaviors, and difficulties empathizing with others often leave them frustrated, misunderstood, teased, and sometimes ostracized.[99]

Researchers have noticed similarities in the characteristics associated with AS and traits stereotypically associated with computer hackers.[100] Tony Atwood, an Australian clinical psychologist, suggests that some computer hackers may have a number of characteristics that are associated with AS.[101] He notes that many diagnosed AS patients are more proficient at computer programming languages than social language, and that the intellectual challenges that are presented by restricted computers and networks may override the illegal nature of their actions. AS has been used as a successful defense in at least one landmark U.K. court case.[102]

Based on over 200 personal interviews with computer criminals, cybercrime expert Donn Parker reports finding significant similarities between AS sufferers and criminal hackers. Many of the computer criminals that Parker interviewed demonstrated the social awkwardness, atypical prosody, and lack of social empathy during social interactions that are characteristic of AS.[103] Anecdotal evidence suggests that computer hackers often have an obsessive interest in technology and computers, similar to that seen in individuals with AS, that forms a salient component of both their individual

and social identities. Due to their egocentric preoccupations, many computer hackers often feel misunderstood and frustrated in face-to-face social situations.

Although the Autism Diagnostic Interview is the most common assessment tool, a number of self-report instruments have been developed to assist in screening disorders on the autism spectrum. To date, sociologist Bernadette Schell has conducted the only empirical study examining Asperger syndrome in a self-identified hacker population.[104] Schell administered the 50-item Autism-Spectrum Quotient (AQ) Inventory to 136 attendees at various well-known hacker conferences (e.g., Defcon[105]) between 2005 and 2007. Previous research using the AQ found that diagnosed AS individuals' mean scores were 35.8 compared to a mean score of 16.4 for a control population. Schell found that the mean score for conference attendees in her study was 19.7. However, 11 percent of males in her sample and 1.5 percent of females had mean AQ scores of 32 or higher. Although this study is limited by due to participants' self-identification as hackers, it does suggest that proposed link between hacker culture and AS may be tenuous at best.[106]

To date there has been no clear empirical evidence to suggest a link between AS and computer crime. There is no evidence whatsoever to suggest that Asperger syndrome causes computer hacking. In fact, most sufferers of AS have been characterized as being extremely honest and lawful citizens.[107] It would be a mistake to assume that all computer hackers are suffering from AS or that every AS sufferer is a computer hacker. Characteristics of AS appear more common in computer *hackers* (i.e., those who explore and tinker with computers and technology), rather than in *crackers* who break into computers or use them for illegal activities. At present, there is still no single all-encompassing personality profile that applies to all computer criminals. In fact, many feel that it is inappropriate to try to create a single personality profile that applies to all computer criminals.

12.5.4 Internet Abuse and Computer Crime. Although *technological addiction* is not a disorder recognized by the American Psychological Association (APA), researchers have suggested that some individuals appear to demonstrate disturbed computer use that is similar to other recognized disorders like compulsive gambling or impulse control disorders. Technological addiction is characterized by:

- Excessive use of a particular technology (usually in reference to computers and Internet usage)
- Preoccupation with the technology
- Systematic increase in use
- Failed attempts to curb one's use
- Feelings of malaise and irritation when not using the technology
- Interference with social and professional pursuits due to excessive technology use[108]

Information security researchers Kent Anderson and Jerrold Post suggest that some computer criminals appear to have symptoms indicative of potentially pathological computer use.[109,110] Research by sociologist Bernadette Schell found that hacker conference participants indicated that they spend on average 24 hours a week on hacker-related activities.[111] Furthermore, Anderson reports that cybercriminals will work for 18 or more hours a day on their computers trying to gain unauthorized access to one

single computer system with little or no external reward for doing so. He also mentions an instance where one U.S. judge even attempted to sentence a computer criminal to psychological treatment for his compulsive computer use.[112]

In their interviews with a number of self-identified computer criminals, sociologists Paul Taylor and Tim Jordan indicated that many of their interviewees report experiencing a thrill or rush that isn't comparable to anything that they experience in their real-world interactions when engaging in illegal activities.[113] A number of their respondents reported feelings of depression, anxiety, and impaired social functioning when they are away from their computers. Taylor and Jordan suggest that these abuse-like characteristics may also be combined with feelings of compulsion regarding computers and new technologies. However, the researchers indicated that these compulsive or obsessive-like characteristics may be as much a function of the information technology (IT) field as they are personality traits.[114] Unlike most disciplines, the IT field is in a constant state of rapid change. To maintain a level of professional expertise and competence in this area, one must devote a good deal of time and resources to monitoring and adapting to this revolutionary field. This need to keep up with the rapidly changing discipline, combined with the euphoric feelings that some experience when committing illegal activities, may increase the likelihood of technological abuse.

Personality theorists state that for some computer criminals, committing electronic crimes produces an experience similar to that of a chemically induced high. Some computer criminals may commit illegal acts because of the euphoric rush they receive from their actions. Information security researcher August Bequai compares the actions of computer criminals to electronic joyriding.[115] Gaining unauthorized access and usage to a computer network is, for these people, similar to that of taking a car on a joyride. One computer cracker interviewed by computer crime researcher Dorothy E. Denning described hacking as "the ultimate cerebral buzz." Other crackers have commented that they received a rush from their illegal activities that felt as if their minds were working at accelerated rates. Some computer criminals have suggested that the euphoric high stems from the dangerous and illegal nature of their activities.[116] Lending empirical support to this idea, criminologist Michael Bachmann found that self-identified computer criminals have a heightened propensity to engage in risk behaviors compared to the general population. Computer criminals have compared the feelings they receive from their illegal intrusions and attacks to the rush that is felt by participating in extreme sports like rock climbing and skydiving.[117]

Researchers have found that some experienced computer users also report sometimes experiencing an altered psychological state known as *flow* while engaging in their technological pursuits. Flow is a psychological state that results in feelings of fulfillment and overall positive affect.[118] When individuals becomes absorbed in a task that matches their skill set, at times they may not be consciously aware of the passage of time or of the differences between the undertaking and their identity. In their study looking at the experience of flow in self-identified computer criminals, psychologists Alexander Voiskounsky and Olga Smyslova found that both inexperienced and highly competent criminals report high levels of flow.[119] For the inexperienced criminals, this flow experience my lead them to limit themselves to low-level challenges, and they may remain in this novice stage for a significant amount of time. More experienced crackers who also experience flow may leave the hacking domain once they are no longer presented with suitable challenges for their abilities. Conversely, they may systematically increase their illegal pursuits in attempts to reacquire the flow state, contributing to technological abuse.[120]

Physical and psychological tolerance can occur when increased amounts of a substance or an activity are needed in order to obtain a *high* or euphoric rush. Tolerance is common in hard drug users who find themselves using increasing amounts of a substance to achieve their original euphoric states. Anecdotal evidence suggests that computer criminals may go through a similar stage of evolution, with each step leading to increased dangerous and riskier behaviors. Many cybercriminals begin by pirating and cracking the copy protection algorithms of software programs. When the *warez* (pirated software) scene loses its thrill, they migrate to chat-room or IRC harassment. The individuals may then begin launching damaging DoS attacks against servers and defacing Websites to obtain that initial rush that originated with simple warez trading. As in a substance abuse, the initial euphoric psychological states and resulting tolerance associated with excessive computer use may explain why some computer criminals repeatedly engage in illicit activities, even after they have been caught and punished.

12.6 ETHICS AND COMPUTER CRIME. Researchers have suggested that computer criminals may have an underdeveloped sense of ethics—a moral immaturity—that contributes to their illegal activities.[121,122,123,124] Because of this ethical immaturity, criminal hackers may think that many of their illegal actions are in fact ethical or beneficial to some degree. Many computer criminals feel that they are ethically entitled to have access to any and all information regardless of legal ownership. Most of these individuals also feel that it is morally right to use inactive computer processing power and time, regardless of who owns the computer system. Computer criminals do not feel that breaking into a computer network should be viewed in the same light as breaking into an individual's house. Often, computer criminals rationalize their illegal activities and justify their behaviors by blaming the victims for not securing their computer networks properly.[125]

Most computer criminals are adolescents, which may account for the underdeveloped sense of ethics in the community.[126] Computer scientist Brian Harvey suggests that due to the relative lack of experience and guidance with the computing environment compared to the real world, teenagers and adolescents may be operating at lower levels of moral functioning when online compared to their interactions and decisions in the real world.[127] Similarly, information security specialist Ira Winkler suggests that computer hackers, because of their generally young age, do not fully understand the repercussions associated with their actions. They also may demonstrate an underdeveloped or complete lack of empathy for their victims.[128] Computer criminals fail to fully realize the consequences of their electronic intrusions into computer networks. The adolescents do not fully comprehend that their mere presence on a computing network could potentially cost companies thousands of dollars, as well as cost systems administrators their jobs.[129]

Sociologist Orly Turgeman-Goldschmidt further suggests that computer criminals may view their behaviors as nothing more than a new form of social entertainment.[130] They see their electronic intrusions as a game that provides them with excitement and thrills. The Internet serves an unlimited playground or social center where "netizens" are able to develop new games and forms of social activities. Computer criminals may view their illegal activity as nothing more than fun and thrill seeking, a new form of entertainment that is carried out on an electronic playground.[131]

According to Winkler, many criminal hackers learn and develop a sense of computing ethics from their online and offline peers.[132] In other words, unlike most instruction on morality and ethics that stems from a responsible adult, computing ethics are socially

learned from other adolescents. Computer crackers learn the rules of hacking and computing from elder statesmen in the hacking community who may be no more than a few years older than themselves. Often, in today's society, the younger generation is more knowledgeable about technology than older adults. When adolescents have problems or need guidance in ambiguous situations, many times their parents or other adult role models are unable to offer them the necessary guidance and assistance. Therefore, the youngsters may seek knowledge from their peers, who may or may not offer them the most ethical or wise advice.[133] In support of this notion, Vincent Sacco and Elia Zureik found that students who viewed illicit computing behaviors as ethical increased the reported likelihood that they would engage in such actions. Computer crime was least reported when the behavior was seen as being more unethical.[134]

While using the Internet, children are constantly making sophisticated judgments without appropriate adult supervision. Adolescents do not have the same ethical maturity that adults have, yet while using the Internet they are given as much power, authority, and responsibility as ethically mature adults.[135] Neil Patrick, the leader of one particularly malicious group of phone-system hackers known as the 414s, stated that he did not know that his hacking was illegal or unethical. In fact, asked when, if ever, he began to question the ethics of his actions, Patrick stated that ethics never came into his mind until the Federal Bureau of Investigation agents were knocking on his front door. Patrick and the other young members of the 414s did not see anything wrong with their actions. To them, breaking into proprietary telecommunications networks was more of a game or challenges rather than a criminal act. They saw nothing ethically wrong with their actions.[136]

Psychologist Lawrence Kohlberg developed a three-level theory to explain normal human moral development.

- The first level deals with avoiding punishments and obtaining rewards.
- The second level emphasizes social rules.
- The third level emphasizes moral principles.

Each of his three levels contains two stages that an individual passes through during adolescence on the way to adult moral development.[137] Computer criminals appear to be operating in the lower three phases of Kohlberg's model: the two stages comprising level 1 and the first stage in level 2. The moral judgments of computer criminals appear to be determined by a need to satisfy their own needs and to avoid disapproval and rejection by others.

In his empirical research on moral development, Rogers found that self-identified computer criminals relied more on hedonistic decision making rather than internal or social morality-based choices, compared to a noncriminal student population. Computer criminals do not appear to be aware of or concerned with the third level of moral development, where moral judgments are motivated by civic respect and one's own moral conscience.[138] Computer criminals may be functioning at the third level of moral development in the physical world, a level appropriate for teens and adults, and may simultaneously be functioning at lower levels of moral development when their online identities are salient.[139] Computer criminals acting at these lower levels of morality may be naively engaging in their illegal activities to satisfy their own curiosity and to gain the approval of their peers, without considering larger moral implications of their behaviors.

According to Shaw and associates, there is a notable lack of ethical regulation and education in organizations, schools, and homes regarding proper computing behavior.[140] Computer criminals who lack ethical maturity fail to realize that their digital actions are sometimes just as damaging as physical aggression. Cybervandals do not see the immediate repercussions of their actions because of the physical distance and lack of social presence in computer-mediated interactions. This ethical immaturity is partially a result of the technology-enhanced knowledge gap between young computer users and their parents. A hacker whose handle was *The Mentor* wrote in the 1986 "Hacker Manifesto,"

> This is our world now … the world of the electron and the switch, the beauty of the baud. We make use of a service already existing without paying for what could be dirt-cheap if it wasn't run by profiteering gluttons, and you call us criminals. We explore … and you call us criminals. We seek after knowledge … and you call us criminals. We exist without skin color, without nationality, without religious bias … and you call us criminals. You build atomic bombs, you wage wars, you murder, cheat, and lie to us and try to make us believe it's for our own good, yet we're the criminals.
>
> Yes, I am a criminal. My crime is that of curiosity. My crime is that of judging people by what they say and think, not what they look like. My crime is that of outsmarting you, something that you will never forgive me for.[141]

In real-world situations, when parents and teachers strive to instill responsible ethics in adolescents, young adults become capable of making informed decisions regarding ethical dilemmas. However, the same adolescents who demonstrate ethical behavior in the physical world may be ethically bereft in cyberspace, partly due to the lack of adult guidance and instruction. Anecdotal evidence suggests that in today's society, adolescents recreationally use, and are more familiar with, computers and the Internet than their parents. These young adults often learn the about Internet-related behaviors and attitudes on their own, or via peer-to-peer interaction. Adolescents are socialized on the Internet by other adolescents, which may lead to a *Lord of the Flies* scenario, where children construct social rules and guidelines to govern their behaviors.[142] These socially constructed norms and guidelines may be both morally and ethically different from real-world norms. Kabay has argued that technological change can take two to three generations for integration of new moral codes into society; by this reasoning, young adults who are growing up with increased awareness of civil behavior in cyberspace will be teaching their own children more appropriate rules of behavior from the earliest ages of the next generation.[143]

12.7 CLASSIFICATIONS OF COMPUTER CRIMINALS. For both ordinary and abnormal behaviors, it is difficult to find one theoretical perspective that can account for every behavior in a given situation. Attitudes and behaviors are the product of the combined influence of an individual's personality and the current social situation. No single theory or theoretical perspective can account for the various types of computerized crimes and the criminals who engage in these activities. There are also many types of computerized crimes, ranging from trading pirated software to cyberterrorism, and many types of computer criminals, ranging from the novice password cracker to the industrial spy. Any theory that would account for the behavior of computer criminals would have to consider, first, the type of illegal activity the person was engaged in and, second, the type of cybercriminal category that the individual falls into.[144] Computer criminals are by nature paranoid and secretive agents who exist in a similar community. They use handles to conceal their true identities and, except for annual hacker

conventions or local meetings, seldom interact with each other in the real world. Therefore, it is difficult for researchers to identify and categorized the various subgroups that exist.

The term *computer hacker* has been both overused and misused as a way of classifying the entire spectrum of computer criminals.[145] The motivations and actions of one subgroup of computer criminals may be entirely different from those of a second group; therefore, it is imperative that in any psychological analysis of computer criminals, the various subcategories be taken into consideration. Many theories have attempted to account for the motivations and behaviors of computer criminals as a whole when the theorists actually were referring to one specific subgroup in the underground culture.[146] Computer criminals are a heterogeneous culture; therefore, one single theory or perspective cannot sufficiently explain all their actions. The fact that researchers traditionally have treated computer criminals as a homogeneous entity has limited the validity and generalizability of their findings. Even researchers who have taken into account the heterogeneous nature of the computing underground have had difficulty with experimental validity. Experimenters have allowed participants to use their own self-classification schemes or attempted to generalize the results of a single subgroup to the entire underground culture.[147]

12.7.1 Early Classification Theories of Computer Criminals. Over the past few decades, several researchers have attempted to develop a categorization system for individuals who engage in various forms of computer crime.[148,149,150,151,152] A comprehensive review of this research is beyond the scope of this chapter. For an extensive review, see Rogers's analysis and development of a new taxonomy for computer criminals.[153] Bill Landreth, a reformed computer cracker, was one of the earliest theorists to develop a classification scheme for computer criminals.[154] His system divided criminals into five categories based on their experience and illegal activities.

1. The *novice* criminals have the least experience with computers and cause the least amount of electronic disruption from their transgressions. They are considered to be tricksters and mischief-makers; for example, AOL users who annoy chatroom members with text floods and DoS-like *punting* programs that crash AOL sessions using specific font or control code strings.
2. The *students* are electronic voyeurs. They spend their time browsing and exploring unauthorized computer systems.
3. The *tourists*, according to Landreth, commit unauthorized intrusions for the emotional rush that results from their actions. This subgroup of computer criminals is similar to Bequai's electronic joyriders.[155] The tourists are thrill-seekers who receive a cerebral buzz from their illegal behaviors.
4. The *crashers* are malicious computer criminals. This subgroup is composed of the darkside criminals that Kabay refers to.[156] The crashers will crack into networks and intentionally delete and destroy data and cause denials of service.
5. Landreth's final classification of computer criminal is the *thieves*. Criminals who fit into this category commit their illegal actions for monetary gain. These individuals are equivalent to the dangerous insiders that Post, Shaw, and Ruby analyzed.[157] Thieves may work alone, or they may be under contract from both foreign and domestic corporations and governments. Examples include the Russian Business Network or RBN.[158]

Former Australian Army intelligence analyst Nicholas Chantler conducted one of the few empirical examinations of computer criminals and their culture.[159] Published in 1995, this survey-based study attempted to gain a deeper understanding of the underground culture as well as to develop a categorization system for cybercriminals. Chantler posted questionnaires to bulletin board systems (BBSs), Usenet newsgroups, and chat rooms owned or frequented by computer criminals. An analysis of the data yielded five primary attributes—criminal activities, hacking prowess, motivations, overall knowledge, and length of time hacking—that Chantler used to create three categories of computer criminals: *lamers*, *neophytes*, and *elites*.[160]

1. *Lamers* have the least technical skill and they have been engaged in their illegal activities for the shortest period of time. This group of criminals is primarily motivated by revenge or theft of services and property.
2. *Neophytes* are more mature than lamers. They are more knowledgeable than the previous category and engage in illegal behaviors in pursuit of increased information.
3. Members of the *elite* group have the highest level of overall knowledge concerning computers and computer crime. They are internally motivated by a desire for knowledge and discovery. They engage in illegal activities for the intellectual challenge and for the thrill they receive from their criminal behaviors.

According to Chantler, the largest proportion of computer criminals at that time, 60 percent, fell into the neophyte category. Thirty percent of computer criminals fell into the elite category, while 10 percent were lamers.[161]

Information security analyst Donn Parker developed a seven-level categorization scheme for computer criminals.[162,163] He formalized his scheme, through years of interaction and structured interviews with computer criminals, into the following categories:

1. **Pranksters** are characterized by their mischievous nature.
2. **Hacksters** are motivated by curiosity and a quest for knowledge. Pranksters and hacksters are the least malicious computer criminals.
3. **Malicious hackers** are motivated by a need for disruption and destruction. They receive pleasure from causing harm to computer systems and financial loss to individuals.
4. **Personal problem solvers** commit illegal activities for personal gain. Problem solvers, the most common type of computer criminal according to Parker, resort to crime after failing in legitimate attempts to resolve their difficulties.
5. **Career criminals** engage in their illegal cyberbehaviors purely for financial gain.
6. **Extreme advocates** have strong ties to religious, political, or social movements. Recently, these types of cybercriminals have been dubbed "hacktivists," a combination of computer hackers and activists.
7. **Malcontents, addicts, and irrational individuals** comprise the final category in Parker's scheme. Individuals in this category usually are suffering from some form of psychological disorder (e.g., antisocial personality disorder).

12.7.2 Rogers's New Taxonomy of Computer Criminals. After an extensive review of past categorization theories, Rogers has advanced an updated continuum of computer criminals based on his previous work.[164] This continuum comprises eight categories based primarily on the criminals' motivations and technological prowess:

1. **Novice (NV)** criminals have the least amount of technical knowledge and skill. Members of this category are relatively new to the scene and use prewritten and precompiled scripts and tools to commit their computerized crimes. They are primarily motivated by the thrill of lawbreaking and making a name for themselves in the underground.
2. **Cyber Punk (CP)** is the group that most fits the traditional stereotype of hacker. Members of this category are slightly more advanced than the novices. These criminals have the ability to create basic attack scripts and programs. Cyber Punks' typical behaviors include Web page defacements, DDoS attacks, carding, and telecommunication fraud. They are motivated by a need for attention, fame, and monetary gain, usually attained by parlaying their crimes into lucrative jobs and book deals. Winkler suggests that the majority of computer criminals fall into either the cyber punk or newbie categories. He estimates that between 35,000 and 50,000 computer criminals, well over 90 percent of their total estimated number, fall into these categories, whom he dubs clueless.[165]
3. **Internals (IN)** consist of disgruntled workers or former workers who hold information technology positions in an organization. Members of this category have an advantage over external attackers due to their job and status within the corporation. Research indicates that internals are responsible for the majority of computer crimes and associated financial loss. Their motivation is typically based on revenge for some perceived wrong (e.g., termination, passed over for a promotion).[166,167] These internal or malicious insiders are examined more fully later in Chapter 13 of this *Handbook*.
4. **Petty Thieves (PT)** are traditional criminals who have turned to technology as a way of keeping up with the times. These individuals are career criminals whose motivation is primarily financial gain from stealing from banks, corporations, and individuals.
5. **Old Guard (OG)** are computer criminals with advanced technical knowledge and skill. These individuals are responsible for writing many of the exploit programs (e.g., stack overflows, rootkits, etc.) that are used by the less knowledgeable novice and cyberpunk crackers in their cyberattacks; however, they are not criminals in the traditional sense. This group has an underdeveloped sense of ethics regarding privacy and intellectual and personal property and engages in behaviors consistent the traditional hacker ethic and ideology described by Levy.[168] Their illegal behaviors are motivated by a quest for knowledge, curiosity, and intellectual stimulation.
6. **Virus Writers (VW)** do not fit neatly into Roger's Taxonomy primarily due to the lack of research on this group of individuals. The demographics and motivations of virus writers are examined more fully later in this chapter.
7. **Professional Criminals (PC)** are traditionally older and more knowledgeable about technology than the previous categories. Members of these categories may be former government and intelligence operatives who are motivated by

financial gain. They may often have access to advanced technology and can be adept at industrial espionage. According to Rogers, PCs may be comprised of ex-intelligence agents and are one of the most dangerous types of computer criminals. Their motivation is primarily large-scale financial gain.

8. **Information Warriors** are highly skilled employees who conduct coordinated attacks against information systems in an attempt to cripple or destabilize the infrastructure. This group may be motivated by allegiance and patriotism. We examine this group more fully in the following section.

Based on these eight categories, Rogers has created a baseline circumplex model of computer criminals that aims to further classify computer criminals. Rogers' circumplex groups computer criminals in a circular diagram using two continua: technological skill and motivation. Using this circumplex, the eight categories can be further subdivided, following future empirical work, to more accurately represent individual subgroups of the computer underground.

12.7.3 Hacktivists and Cyberterrorists: Hacking for a Cause. With the emerging protests against the economically advantaged 1 percent and the political upheaval in Middle East, new trends in cybercrime have emerged in the form of a revitalized hacktivist ethos. *Hacktivism* is defined as cause-based hacking for social, political, patriotic, or religious purposes. In general, the perpetrators engage in technological dissent in order to promote freedom of speech and fair distribution of wealth and to combat censorship. Hactivists will often rally behind a symbol, logo, or flag, whether they are religious, political, or emblematic. The term was first used by groups like Electronic Disturbance Theatre (EDT) and the Cult of the Dead Cow (cDc) in the mid-1990s to establish digital protests and *hack-ins* where the groups would launch coordinated DoS attacks and mass Web page defacements as a means of online protest.[169,170,171,172] Paralleling the global "Occupy" protest movements, many of these modern hacktivist groups have no central leadership or core philosophy, making their actions and targets difficult to predict. They present an increased danger due to their tendency to single out and target high-profile people and organizations to achieve maximal exposure for their group, cause, and interests. Hacktivists are not motivated by financial gain, but instead prefer to publicize their cause through the harm and embarrassment of their victims.[173] The methods employed by modern hacktivists have evolved from simple DoS attacks and Web defacements to massive amounts of government, corporate, and personal data theft. The number of hacktivist attacks in 2011 surpassed all of the previous year's attacks combined that were attributed to social and political motivations.[174]

According to Rogers, *cyber-terrorists*, along with *professional criminals*, may present the most danger to individuals and organizations. A cyber-terrorist is defined as, "an individual who uses computer or network technology to control, dominate, or coerce through the use of terror in furtherance of political or social objectives."[175] Following the September 11, 2001, attacks in the United States, the term cyberterrorism has been frequently overused and misused by the media. An individual who stumbles upon a vulnerable .mil or .gov Website and decides to deface it should not be considered a cyberterrorist unless there is a premeditated intent to cause panic and fear, with the object of obtaining some social end.

Although most of the Web page defacements and network attacks labeled by the media as cyberterrorism would not fit Rogers's definition, that is not to say that the Internet will not be used as means for terrorist acts in the future.[176] Spurred on by

the specter of cyberterrorism and claims by hackers that they could cripple the Internet in 30 minutes, governments and corporations are investing millions into research aimed at protecting the global computing infrastructure from cyberattacks. Rogers suggests that the relative anonymity of the perpetrator, and the multitude of potential targets that could be simultaneously attacked, makes the Internet a very appealing target for terrorists' actions that will likely be exploited in the near future.[177]

Perhaps the most infamous case of cyberterrorism occurred in 2010 when researchers analyzed Stuxnet, one of the most sophisticated electronic worms ever discovered in the wild. Stuxnet was created with the singular goal of creating physical damage to centrifuges in one of Iran's nuclear enrichment facilities. Although no one has taken credit for the Stuxnet worm, computer security researchers suggests that it may have been cyberwarriors working for the United States and/or Israeli governments.[178] Theorists suggest that new breeds of computer criminals, dubbed cyberterrorists and hacktivists, are emerging and are motivated by political or social ideologies related information freedom, nationalism, ethnic pride, and warfare.[179,180]

12.7.4 Dangerous/Malicious Insiders (DI/MI). Dr. Eric Shaw and associates classify computer criminals into two categories: outside intruders and dangerous insiders.[181,182,183] The researchers focus on the critical IT insiders who are typically programmers, technical support staff, networking operators, administrators, consultants, and temporary workers in organizations. Malicious insiders (MI) are a subgroup of such employees who are motivated by greed, revenge, problem resolution, and ego gratification. Shaw and colleagues estimate that the theft of trade secrets costs more than $250 billion annually for U.S. businesses alone.

Shaw's research, based on corporate surveys and hundreds of investigations by the U.S. Secret Service and Carnegie Mellon University's Computer Emergency Response Team Coordination Center (CERT-CC), has attempted to compile an initial profile of dangerous insiders. Their critical-pathway approach identifies psychological and situational precursors that may predispose an employee to engage in insider espionage[184]:

1. **Personal Predisposition**: medical/psychiatric problems, reduced social skills, previous violations of the law or business ethics, social or professional ties with competitors or adversaries
2. **Personal Life Stressors**: financial burden, relationship problems, medical issues, legal issues
3. **Professional Stressors**: demotions, failed promotions, poor performance review, transfer, supervisor disagreements, looming layoffs
4. **Concerning Behaviors**: workplace violations, conflict, intellectual property (IP) disagreements
5. **Maladaptive Organizational Responses**: failure to detect, investigate, appreciate, and appropriately respond to concerning behaviors of employees

The generic profile for an MI is a male in his late thirties who has a job in a technology-related area within the company. The majority of thefts occur when an MI has accepted a job elsewhere and within a month of starting their new job. Moore, as cited by Shaw, classifies MIs into two categories, the *Entitled Disgruntled Thief* and the *Machiavellian Leader*.

- The Entitled Disgruntled Thief typically engages in IP theft due to professional stressors. They steal information that they directly worked on or helped to develop. These individuals feel entitled to the information and typically will use it to help them acquire a new job or to enhance their performance at their already acquired position.
- The actions of the Machiavellian Leader are more premeditated. They are driven more by personal ambition rather than personal or professional stressors. These individuals may also recruit coworkers to assist them in their thefts.

According to Shaw and coauthors, these dangerous insiders typically have introverted personalities.[185] They demonstrate a preference for solitary intellectual activities over interpersonal interaction. Members of this subgroup may have had numerous personal and social frustrations that have hindered their interpersonal interactions and contributed to their antiauthoritarian attitudes.

The malicious subgroup of critical information technologists has been characterized as having a loose or immature sense of ethics regarding computers. The dangerous insiders rationalize their crimes by blaming their company or supervisors for bringing any negative consequences on themselves. They feel that any electronic damage they cause is the fault of the organization for treating them unfairly. The researchers also note that many insiders identify more with their profession than with the company for which they work. This heightened identification with the profession undermines an insider's loyalty to an organization. This reduced loyalty is evidenced by the high turnover rates of jobs in the IT industry. According to Shaw and associates, the unstable bond between insiders and their organizations creates undue tension with regard to security practices and IP rights.[186]

Researchers also have suggested that dangerous insiders are characterized by an increased sense of entitlement and hostility toward organizational authority. According to Shaw and coauthors, when an unfulfilled sense of entitlement is combined with previous hostility toward authority, malicious acts or revenge against the organization are typical.[187,188,189]

12.7.5 Virus Creators. Unlike more traditional forms of computer crime, there has been very little research examining the motivations and behaviors of malware writers, usually referred to as *virus writers*.[190] Despite the name, the group writes other forms of malware such as worms and Trojans. The limited number of research reports on this particular subgroup of computer criminal has relied primarily on one-on-one interviews and surveys, in an effort to understand the actions and motivations of virus creators.[191,192,193,194]

Using case studies and multiple interviews, researchers Andrew Bissett and Geraldine Shipton examined the factors that influence and motivate virus writers.[195] The researchers suggest that it is difficult to generalize their findings to all virus creators because of the limited published literature and research regarding virus writers. Virus creators appear to demonstrate conscious motivations for their potentially destructive actions that are similar to the motivations of traditional computer criminals. The coders create and distribute their software for reasons of nonspecific malice, employee revenge, ideological motives, commercial sabotage, and information warfare.

Bissett and Shipton's review of anti-virus expert Sarah Gordon's interview with Dark Avenger reveals some of the motivations behind one of the most notorious virus writers.[196] They suggest that Dark Avenger consistently denies responsibility for his

creations and, like traditional computer criminals, engages in victim blaming. Dark Avenger states that it is human stupidity, not the computer, that spreads viruses. The virus writer also appears to self-identify with his malicious code. Dark Avenger seems to project his persona onto viruses in a process called projective identification. During the interview, Dark Avenger stated that the United States could prevent him from entering the country, but it is unable to stop his viruses. Dark Avenger also attempted to justify creating destructive viruses by commenting that most personal computers did not store data of any value, and therefore his malicious programs were not doing any real harm. Similar to the motivations spurring dangerous insiders, the researchers suggest that Dark Avenger creates malicious viruses because he is envious of the work and achievements of other computing professionals. In the interview, Dark Avenger commented that he hates it when people have more powerful computing resources than he does, especially when the individuals do not use the resources for anything that he deems constructive.[197]

Sarah Gordon has also examined the ethical development of several virus writers using surveys and structured interviews.[198,199,200] Her initial four case studies involved an adolescent virus writer, a college-age virus writer, a professionally employed virus writer, and an ex–virus writer. The interviews revealed that all four individuals appeared to demonstrate normal ethical maturity and development consistent with Kohlberg's previously reviewed stage theory.[201] Gordon suggests that there appear to be many different reasons why individuals create and distribute viruses, including boredom, exploration, recognition, peer pressure, and sheer malice.[202]

Gordon suggests that the virus underground in the mid and late 1990s was populated by a second generation or next generation of virus writers whose skill and ability at virus construction is comparable to that of the old school, original virus writers.[203] On the surface, these second-generation creators maintain a public façade that suggests that they are extremely cruel, obnoxious, and more technologically advanced than the previous generation. The next-generation virus writers appear to be more aware of the ethical responsibilities surrounding virus creation and distribution; however, the exact definition of "responsible virus creation and distribution" varies from individual to individual. Many of these next-generation virus writers have considerable technical skill and are motivated by the challenge to defeat the malware countermeasures implemented by antivirus vendors.[204]

According to Gordon, another group of virus creators populating the virus underground in the 1990s was composed of "new-age" virus writers.[205] These individuals are motivated by current trends, such as political activism, virus exchange, freedom of information, and challenges to write the most destructive or sophisticated virus, as opposed to technical exploration. These virus writers are motivated by boredom, intellectual curiosity, mixed messages surrounding the legality of virus creation, and increased access to ever-more-powerful technological resources. Gordon suggests that these new-age virus writers may be older and wiser than the second or next-generation creators. They are very selective as to who has access to their creations, and they do not share their findings or accomplishments with members outside their respective group. Unlike the next-generation creators, new-age virus writers will not stop or grow out of writing viruses, as they are most likely already adults. They will continue to write and distribute more sophisticated viruses in part due to the mixed messages concerning the ethical nature of virus creation propagated by the popular media, academia, and popular freedom of information zeitgiest.[206]

Researchers have stated that we must be careful not to view virus creators as a homogeneous group.[207,208] Instead, we must monitor the virus-exchange community

while pursuing in-depth case histories that may aid in our understanding of virus writers. Education about the ethical nature of virus creation and distribution, and about the repercussions associated with malicious code, may attenuate these potentially destructive activities.

12.8 RECOMMENDATIONS. Psychological theories offer various explanations that may influence criminal activities on the Internet. These situational influences may interact with various individual personality traits to further contribute to illegal behaviors. The current task for researchers is to untangle these personality and situation influences on electronic behavior. They must determine what situations and characteristics are influencing the various subgroups of computer criminals. There is no simple explanation as to why computer criminals engage in hostile and destructive acts. The answer lies in a complex mixture of factors that depends on the social environment and individual personality factors. There are numerous types of computer criminals ranging from script kids to the professional criminal, each with varying personalities and motivators. The interaction of personality variables and environmental factors will determine how a computer criminal reacts in any given situation.

One underlying theme in reducing the overall prevalence of computer crime has been to remove the tangible and psychological reward system that currently surrounds the culture of the Internet underground. Criminal law researchers A. S. Dalal and Raghav Sharma suggest that the information security field has established a pattern of rewarding criminal hackers for their exploits by offering them employment opportunities that further undermines law enforcement efforts to combat computer crime.[209] Preventing cybercriminals from profiting from their transgressions via lucrative jobs and book deals may result in a socially learned deterrent to engaging in cybercrime. Thomas Holt and colleagues found strong support for the social learning model of cybercrime in their research: those who received reinforcement in the form of encouragement, praise, and resources from peers and authority figures were more likely to engage in cybercrime activities. This model suggests that removing this reinforcement would significantly reduce the amount and frequency of cyberdeviance, especially in the case of criminals that fall into Rogers' *novice* and *cyberpunk* categories.[210,211]

Cautioning parents, teachers, and bosses about the dangers of praising adolescents for cyberdeviance and instead instituting consistent punishments should also help to mitigate computer crime for novices. Also, a more proactive strategy that parents and educators could take would be to engage in moral or ethics training on the use of computers and technology before children begin independently using technology on a regular basis. Then, when confronted with ambiguous decisions while using the computers, children would reflexively rely on adults' moral advice and teachings instead of their peers'. Scholastic Incorporated demonstrated in a recent survey that 48 percent of elementary and middle school students did not consider computer hacking to be a crime.[212] Recently, the U.S. Department of Justice (DoJ) formed a cybercitizen awareness program in an effort to educate parents, teachers, and children about ethical and unethical computing practices. The program seeks to educate individuals through ethics conferences, multimedia presentations, and speaking engagements at schools around the country. This program and others like it may aid to increase moral responsibility and ethical behaviors of adolescents on the Internet.[213]

Dalal and Sharma also suggest that punishments take into account the motivation and type of cybercrime that is committed, instead of using a rigid approach to sentencing that precludes adapting punishment to the nature of the criminal. For instance, longer jail terms may be a suitable deterrent for criminals who are motivated by financial

gain or malice (e.g., internals, petty thieves, and malicious insiders), but might not be warranted for those who are motivated by intellectual challenges and curiosity (e.g., old guard hackers).

Another problem may be the relatively light sentences that are typically handed down for computer crimes, and the perpetrators' perceptions that they are unlikely to be caught or prosecuted for their transgressions. One study found that the perceived gains achieved by illegal hacking far outweighed the potential costs associated with digital crime. Furthermore, even though computer criminals feel that the punishments for their crimes are severe enough, their feelings that there is a very low probability that they'll be caught minimizes the effects of punishment severity as a deterrent.[214] Even still, the notoriety that a computer criminal gets from being handed a jail term can be parlayed into a lucrative job once the individual has paid the price for their criminal endeavors.

A more proactive approach to reducing the motivation to engage in illicit activities while using computers could be to establish legal hacking networks for novices to experiment with and explore. For novices and even cyberpunks, safe hacking spaces where they have unlimited access to networks and programs may reduce the psychological reactance that results from strictly enforced computing access and resources. These legal hacking networks would also serve to redirect the users' focus on the technical aspect of computing and traditional hacking endeavors rather than trying to crack open the system. By redirecting adolescents' curiosity toward traditional hacking and technical pursuits into an open network, the users may be less inclined to engage in criminal pursuits.

To combat the threat of malicious insiders, Shaw suggests that businesses also adopt a proactive approach to threat mitigation.[215] Developing more comprehensive employee-screening methods that include social-media reviews, background checks, substance abuse tests, honesty/ethical tests, and psychological screenings may serve to mitigate the risks of IP theft. Additionally, a more cost-effective and time-efficient method of reducing IP theft would be increasing employee awareness of IP nondisclosure agreements (NDAs) and ensuring that they are aware of the consequences of violating these agreements. Increasing employees' and managers' awareness of intellectual property rights and risk factors, as well as instituting a systematic reporting system for potential violation could mitigate some of the risks of IP theft. Shaw reports that the majority of IP theft was detected by nontechnical means (e.g., employees noticing suspicious activity or similar products being marketed by competitors).[216]

This chapter summarized research on the psychological motivations of computer criminals. The cybercrime landscape is vast in scope and populated by subgroups of computer criminals with their own patterns of motivations, goals, attitudes, and behaviors. Psychological theory and research is one step in trying to generate effective countermeasures to combat the problem of cybercrime.

Although there is no shortage of theories and anecdotal evidence to account for why cybercriminals engage in network intrusions, there is a marked lack of empirical evidence that serves to test and validate these hypotheses. Information security professionals, criminologists, and psychologists must find ways to begin jointly developing, testing, and examining these ideas in order to find concrete solutions for the problem of cybercrime. Such effort will be facilitated by government and private-sector funding to support research into the psychological and social dynamics of the Internet underground. The information security industry as a whole would do well to support these applied scientific efforts.[217]

12.9 FURTHER READING

Jaishankar, K., ed. *Cyber Criminology: Exploring Internet Crimes and Criminal Behavior*. Boca Raton, FL: CRC Press, 2011.

Kirwan, G., and A. Power, eds. *The Psychology of Cyber Crime: Concepts and Principles*. *Hershey, PA*: IGI Global, 2011.

12.10 NOTES

1. Symantec, *Internet Security Threat Report: 2011 Trends, Vol. 17*, Symantec Website, April 2012, www.symanteccloud.com/en/us/mlireport/b-istr_main_report_2011_21239364.en-us.pdf
2. Marc Rogers, "The Development of a Meaningful Hacker Taxonomy: A Two Dimensional Approach," *CERIAS Technical Report*, 43 (2005): 1–9.
3. National Institute of Justice, "Electronic Crime Research and Development," 2006, www.nij.gov/topics/crime/internet-electronic/
4. Symantec, *Internet Security*, 2012.
5. For more information about computer crimes, see Chapter 2 in this Handbook.
6. Verizon, "Data Breach Investigations Report," Verizon Website, 2012, www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf
7. Symantec, *Internet Security*, 2012.
8. For more information about information warfare, see Chapter 14 in this *Handbook*.
9. Douglas Campbell, "A Detailed History of Terrorist and Hostile Intelligence Attacks Against Computer Resources," 1992. Available at www.syneca.com/publications.htm
10. Donn Parker, "How to Solve the Hacker Problem," *Journal of the National Computer Security Association*, No. 5 (1994), pp. 4–8.
11. Jerrold Post, "The Dangerous Information Systems Insider: Psychological Perspectives," (1998). Available email: jmpost@pol-psych.com
12. Parker, "Hacker Problem," 1994 (pp. 4–8).
13. Paul Taylor, *Hackers: Crime in the Digital Sublime*, (London & New York: Routledge, 1999).
14. Raoul Chiesa, Stefania Ducci, and Silvio Ciappi, *Profiling Hackers: The Science of Criminal Profiling as Applied to the World of Hacking* (Boca Raton, FL: Auerbach Publications, 2008).
15. Orly Turgeman-Goldschmidt, "Hackers' Accounts: Hacking as a Social Entertainment," *Social Science Computer Review*, 23, no. 1 (2005): 8–23.
16. Tim Jordan and Paul Taylor, "A Sociology of Hackers," *Sociological Review*, 46, no. 4 (1998): 757–780.
17. Taylor, *Hackers*, 1999.
18. Turgeman-Goldschmidt, "Hackers' Accounts," 2005.
19. Taylor, *Hackers*, 1999.
20. Taylor, *Hackers*, 1999.
21. Chiesa, Ducci, and Ciappi, *Profiling Hackers*, 2008.
22. Chiesa, Ducci, and Ciappi, *Profiling Hackers*, 2008.

23. Chiesa, Ducci, and Ciappi, *Profiling Hackers*, 2008.
24. Jordan and Taylor, "A Sociology," 1998.
25. Hyung-jin Woo, Yeora Kim, & Joseph Dominick, "Hackers: Militants or Merry Pranksters? A Content Analysis of Defaced Web Pages," *Media Psychology*, 6, no. 1 (2004): 63–83.
26. David G. Myers, *Social Psychology*, 8th ed. (New York: McGraw-Hill, 2005).
27. Myers, *Social Psychology*, 2005.
28. Woo, Kim, and Dominick, "Hackers: Militants or Merry Pranksters?" 2004.
29. For more information on denial-of-service attacks, see Chapter 18 in this Handbook.
30. Stanley Milgram, *Obedience to Authority* (New York: Harper & Row, 1974).
31. Albert Bandura, "Selective Activation and Disengagement of Moral Control," *Journal of Social Issues*, 46, (1990): 27–46.
32. Marc Rogers, "Modern-Day Robin Hood or Moral Disengagement?" 1999. http://homes.cerias.purdue.edu/~mkr/moral_doc.pdf
33. Turgeman-Goldschmidt, "Hackers' Accounts," 2005.
34. Jordan and Taylor, "A Sociology," 1998.
35. Woo, Kim, & Dominick, "Hackers: Militants or Merry Pranksters?" 2004.
36. Rogers, "Modern-Day Robin Hood," 1999.
37. Randall Young, Lixuan Zhang, and Victor R. Prybutok, "Hacking into the Minds of Hackers," *Information Systems Management*, 24 (2007): 281–187.
38. Sara Kiesler, Jane Siegel, and Timothy McGuire, "Social Psychological Aspects of Computer-mediated Communication," *American Psychologist*, 39 (1984): 1123–1134.
39. Sara Kiesler and Lee Sproull, "Group Decision Making and Communication Technology," *Organizational Behavior and Human Decision Processes*, 52 (1992): 96–123.
40. Kiesler and Sproull, "Group Decision Making," 1992.
41. Kiesler, Siegel, and McGuire, "Social Psychological Aspects," 1984.
42. Kiesler and Sproull, "Group Decision Making," 1992.
43. Turgeman-Goldschmidt, "Hackers' Accounts," 2005.
44. Phillip Zimbardo, "The Human Choice: Individuation, Reason, and Order Versus Deindividuation, Impulse, and Chaos," Nebraska Symposium on Motivation, No. 17 (1969): 237–307.
45. Edward Diener, "Deindividuation, Self-Awareness, and Disinhibition," *Journal of Personality and Social Psychology*, 37 (1979): 1160–1171.
46. Zimbardo, "The Human Choice," 1969.
47. Kiesler, Siegel, and McGuire, "Social Psychological Aspects," 1984.
48. Zimbardo, "The Human Choice," 1969.
49. Zimbardo, "The Human Choice," 1969.
50. M. E. Kabay, "Anonymity and Pseudonymity in Cyberspace: Deindividuation, Incivility and Lawlessness versus Freedom and Privacy," Annual Conference of the European Institute for Computer Anti-virus Research (EICAR), 1998, revised 2001, www.mekabay.com/overviews/anonpseudo.pdf

51. Tom Postmes, Russell Spears, and Martin Lea, "Social Identity, Normative Content, and 'Deindividuation' in Computer-Mediated Groups," In *Social Identity: Context, Commitment, Content*, ed. N. Ellemers, R. Spears, and B. Doosje (Oxford: Blackwell, 1999), 164–183.
52. Tom Postmes, Russell Spears, and Martin Lea, "Breaching or Building Social Boundaries? SIDE-Effects of Computer Mediated Communication," *Communication Research*, 25 (1998): 689–715.
53. Stephen D. Reicher, Russell Spears, and Tom Postmes, "A Social Identity Model of Deindividuation Phenomena," *European Review of Social Psychology*, 6 (1995): 161–198.
54. Henri Tajfel and John C. Turner, "The Social Identity Theory of Inter-Group Behavior," In *Psychology of Intergroup Relations*, ed. S. Worchel and L. W. Austin (Chicago: Nelson-Hall, 1986), 2–24.
55. Postmes, Spears, and Lea, "Social Identity," 1999.
56. Postmes, Spears, and Lea, "Breaching or Building," 1998.
57. Reicher, Spears, and Postmes, "A Social Identity Model," 1995.
58. Jordan and Taylor, "A Sociology," 1998.
59. Tajfel and Turner, "The Social Identity Theory," 1986.
60. Woo, Kim, & Dominick, "Hackers: Militants or Merry Pranksters?" 2004.
61. Woo, Kim, & Dominick, "Hackers: Militants or Merry Pranksters?" 2004.
62. Woo, Kim, & Dominick, "Hackers: Militants or Merry Pranksters?" 2004.
63. Woo, Kim, & Dominick, "Hackers: Militants or Merry Pranksters?" 2004.
64. Rogers, "Modern-Day Robin Hood," 1999.
65. Albert Bandura, "The Social Learning Perspective: Mechanisms of Aggression." In *Psychology of Crime and Criminal Justice*, ed. H. Toch (New York: Holt, Rinehart & Winston, 1979).
66. Bandura, "Social Learning Perspective," 1979.
67. Thomas J. Holt, George W. Burruss, & Adam M. Bossler, "Social Learning and Cyber-Deviance: Examining the Importance of a Full Social Learning Model in the Virtual World," *Journal of Crime & Justice*, 33, no. 2 (2010): 31–61.
68. M. E. Kabay, "Totem and Taboo in Cyberspace: Integrating Cyberspace into Our Moral Universe," *Journal of the National Computer Security Association* (1996): 4–9. www.mekabay.com/ethics/totem_taboo_cyber.pdf
69. Chiesa, Ducci, and Ciappi, *Profiling Hackers*, 2008.
70. Kabay, "Totem and Taboo," 1996.
71. Holt, Burruss, and Bossler, "Social Learning and Cyber-Deviance," 2010.
72. Bandura, "Social Learning Perspective," 1979.
73. Myers, *Social Psychology*, 2005.
74. American Psychiatric Association, *Diagnostic and Statistical Manual of Mental Disorders*, 4th ed. (DSM-IV) (Washington, DC: American Psychiatric Association, 1994).
75. American Psychiatric Association, *DSM-IV*, 1994.
76. Kabay, "Totem and Taboo," 1996.
77. Eric D. Shaw, Keven Ruby, and Jerrold Post, "The Insider Threat to Information Systems: The Psychology of the Dangerous Insider," *Security Awareness Bulletin*, 2 (1998): 1–10.

78. Kabay, "Totem and Taboo," 1996.
79. Shaw, Ruby, and Post, "Insider Threat," 1998.
80. American Psychiatric Association, *DSM-IV*, 1994.
81. Shaw, Ruby, and Post, "Insider Threat," 1998.
82. Toxic Shock, "Another View of Hacking: The Evil That Hackers Do," *Computer Underground Digest*, 2 (1990). Available: http://cu-digest.org/CUDS2/cud2.html
83. Shaw, Ruby, and Post, "Insider Threat," 1998.
84. Toxic Shock, "Another View," 1990.
85. Marc Rogers, "Understanding Criminal Computer Behavior: A Personality Trait and Moral Choice Analysis," 2003, Marc Rogers's Website, http://homes.cerias.purdue.edu/~mkr/CPA.doc
86. Robert R. McCrae and Paul T. Costa, Jr. "Personality Trait Structure as a Human Universal," *American Psychologist*, 52. (1997): 509–516.
87. Rogers, "Criminal Computer Behavior," 2003.
88. Marcus Rogers, Kathryn Seigfried, and Kirti Tidke, "Self-Reported Computer Criminal Behavior: A Psychological Analysis," *Digital Investigation,* 3S (2006): 116–120.
89. Postmes, Spears, & Lea, "Social Identity," 1999
90. Sherry Turkle, "Identity Crisis," in *Life on the Screen: Identity in the Age of the Internet* (New York: Simon & Schuster, 1995): 255–269.
91. M. J. Zuckerman, "Hacker Reminds Some of Asperger Syndrome," *USA Today*, March 3, 2001, www.usatoday.com/news/health/2001-03-29-asperger.htm
92. Suelette Dreyfus, "Cracking the Hackers' Code," *The Sydney Morning Herald*, August 8, 2002, http://smh.com.au/articles/2002/08/20/1029114072039.html
93. Bernadette H. Schell and June Meluychuk, "Female and Male Hacker Conferences Attendees: Their Autism-Spectrum Quotient (AQ) Scores and Self-Reported Adulthood Experiences," in *Corporate Hacking and Technology-Driven Crime: Social Dynamics and Implications*, ed. Thomas J. Holt and Bernadette H. Schell (Hershey, PA: IGI Global, 2011): 144–166.
94. American Psychiatric Association, *DSM-IV*, 1994.
95. American Psychiatric Association, *DSM-IV*, 1994.
96. Stephen Bauer, "Asperger Syndrome," 2001, http://aspergersyndrome.org/
97. Bauer, "Asperger Syndrome," 2001.
98. Peter Smith, "The Cybercitizen Partnership: Teaching Children Cyber Ethics," Cybercitizen Partnership, 2000, www.cybercitizenship.org/ethics/whitepaper.html
99. Bauer, "Asperger Syndrome," 2001.
100. Zuckerman, "Hacker Reminds Some," 2001.
101. Dreyfus, "Cracking the Hackers' Code," 2002.
102. Dreyfus, "Cracking the Hackers' Code," 2002.
103. Zuckerman, "Hacker Reminds Some," 2001.
104. Schell and Meluychuk, "Female and Male Hacker AQ Scores," 2011.
105. DEFCON Website, www.defcon.org
106. Schell and Meluychuk, "Female and Male Hacker AQ Scores," 2011.

107. Bauer, "Asperger Syndrome," 2001.
108. Mark Griffiths, "Internet Addiction: Does It Really Exist?" In *Psychology and the Internet: Intrapersonal, Interpersonal and Transpersonal Applications*, ed. J. Gackenbach, (New York: Academic Press, 1998): 61–75.
109. K. E. Anderson, "International Intrusion: Motives and Patterns," 1994, www.aracnet.com/~kea/Papers/paper.shtml
110. Jerrold Post, Eric Shaw, and Keven Ruby, "Information Terrorism and the Dangerous Insider," Paper presented at the InfowarCon, 1998, Washington, D.C.
111. Schell and Meluychuk, "Female and Male Hacker AQ Scores," 2011.
112. Schell and Meluychuk, "Female and Male Hacker AQ Scores," 2011.
113. Jordan and Taylor, "A Sociology," 1998.
114. Jordan and Taylor, "A Sociology," 1998.
115. August Bequai, *Technocrimes* (Lexington, MA: Lexington Books, 1987).
116. Dorothy E. Denning, "Concerning Hackers Who Break into Computer Systems," *CPSR* (1990), http://cyber.eserver.org/hackers.txt
117. Michael Bachmann, "The Risk Propensity and Rationality of Computer Hackers," *International Journal of Cyber Criminology*, 4 (2010): 643–656.
118. Csikszentmihalyi, Mihaly, *Flow* (Harper Perennial Modern Classics, 2008).
119. Alexander E. Voiskounsky and Olga V. Smyslova, "Flow-Based Model of Computer Hackers' Motivation," *CyberPsychology & Behavior*, 6, no. 2 (2003): 171–161.
120. Voiskounsky & Smyslova, "Flow-Based Model," 2003.
121. Rogers, "Modern-Day Robin Hood," 1999.
122. Post, Shaw, and Ruby, "Information Terrorism," 1998.
123. Denning, "Concerning Hackers," 1990.
124. Ira Winkler, "Why Hackers Do the Things They Do?" *Journal of the National Computer Security Association*, 7 (1996): 12.
125. Woo, Kim, and Dominick, "Hackers: Militants or Merry Pranksters?" 2004.
126. Denning, "Concerning Hackers," 1990.
127. Brian Harvey, "Computer Hacking and Ethics," Brian Harvey's Website, 1998, www.cs.berkeley.edu/~bh/hackers.html
128. Winkler, "Why Hackers Do," 1996.
129. Harvey, "Computer Hacking and Ethics," 1998.
130. Turgeman-Goldschmidt, "Hackers' Accounts," 2005.
131. Turgeman-Goldschmidt, "Hackers' Accounts," 2005.
132. Winkler, "Why Hackers Do," 1996.
133. Winkler, "Why Hackers Do," 1996.
134. Vincent Sacco and Elia Zureik, "Correlates of Computer Misuse: Data from a Self-Reporting Sample," *Behavior & Information Technology*, 9 (1990): 353–369.
135. Harvey, "Computer Hacking and Ethics," 1998.
136. Harvey, "Computer Hacking and Ethics," 1998.
137. Myers, *Social Psychology*, 2005.
138. Rogers, "Modern-Day Robin Hood," 1999.

139. Harvey, "Computer Hacking and Ethics," 1998.
140. Shaw, Ruby, and Post, "Insider Threat," 1998.
141. Mentor, "The Hacker Manifesto," *Phrack Magazine*, January 8, 1986, www.mithral.com/~beberg/manifesto.html
142. William Golding, *Lord of the Flies* (London: Faber and Faber, 1954).
143. Kabay, "Totem and Taboo," 1996.
144. Denning, "Concerning Hackers," 1990.
145. Marc Rogers, "A New Hacker Taxonomy," Marc Rogers's Website, 2000, http://homes.cerias.purdue.edu/~mkr/hacker_doc.pdf
146. Rogers, "A New Hacker Taxonomy," 2000.
147. Rogers, "A New Hacker Taxonomy," 2000.
148. Rogers, "Meaningful Hacker Taxonomy," 2005.
149. A. N. Chantler, "Risk: The Profile of the Computer Hacker." Ph.D. dissertation, Curtin University of Technology, 1995.
150. Rogers, "A New Hacker Taxonomy," 2000
151. Bill Landreth, *Out of the Inner Circle* (Redmond, WA: Microsoft Books, 1985).
152. Donn Parker, *Fighting Computer Crime: A New Framework for Protecting Information* (New York: John Wiley & Sons, 1998).
153. Rogers, "Meaningful Hacker Taxonomy," 2005.
154. Landreth, *Out of the Inner Circle*, 1985.
155. Bequai, *Technocrimes*, 1987.
156. Kabay, "Totem and Taboo," 1996.
157. Post, Shaw, and Ruby, "Information Terrorism," 1998.
158. See B. Guinen and M. E. Kabay, "The Russian Cybermafia: Beginnings," 2011, www.mekabay.com/nwss/866_russian_cybercrime_(guinen)_part_1.pdf; M. E. Kabay and B. Guinen, "The Russian Cybermafia: Boa Factory & CarderPlanet," 2011, www.mekabay.com/nwss/867_russian_cybercrime_(guinen)_part_2.pdf; and B. Guinenand M. E. Kabay, "The Russian Cybermafia: RBN & the RBS WorldPay Attack," 2011, www.mekabay.com/nwss/868_russian_cyber crime_(guinen)_part_3.pdf
159. Chantler, "Risk," 1995.
160. Chantler, "Risk," 1995.
161. Chantler, "Risk," 1995.
162. Rogers, "A New Hacker Taxonomy," 2000.
163. Parker, *Fighting Computer Crime*, 1998.
164. Rogers, "Meaningful Hacker Taxonomy," 2005.
165. Winkler, "Why Hackers Do," 1996.
166. Eric D. Shaw, "The Role of Behavioral Research and Profiling in Malicious Cyber Insider Investigations," *Digital Investigation*, 3 (2006): 20–31.
167. Eric D. Shaw and Harley V. Stock, "Behavioral Risk Indicators of Malicious Insider Theft of Intellectual Property: Misreading the Writing on the Wall," Symantec White Paper, 2011, https://www4.symantec.com/mktginfo/whitepaper/21220067_GA_WP_Malicious_Insider_12_11_dai81510_cta56681.pdf
168. Steven Levy, *Hackers* (New York: Dell Publishing, 1984).

169. Samantha Murphy, "Culture Lab: Inside the Hacktivists Revolution," *Culture-Lab | New Scientist*, April 12, 2012, www.newscientist.com/blogs/culturelab/2012/04/samantha-murphy-contributorit-is-the.html
170. Thomas J. Holt, "The Attack Dynamics of Political and Religiously Motivated Hackers," In *Cyber Infrastructure Protection*, ed. Tarek Saadawi and Louis Jordan Jr., (Strategic Studies Institute, 2011): 159–180, www.strategicstudiesinstitute.army.mil/pdffiles/PUB1067.pdf
171. Marc Rogers, "The Psychology of Cyber-Terrorism," In *Terrorists, Victims, and Society: Psychological Perspectives on Terrorism and its Consequences*, ed. A. Silke (London: Wiley & Sons, 2003).
172. Alexander Gostev and Costin Raiu, "Kaspersky Security Bulletin. Malware Evolution 2011," Securelist Website, 2012, www.securelist.com/en/analysis/204792217/Kaspersky_Security_Bulletin_ Malware_Evolution_2011
173. Holt, "Attack Dynamics," 2011.
174. Verizon, "Data Breach Investigations Report," 2012.
175. Rogers, "Psychology of Cyber-Terrorism," 2003.
176. Rogers, "Psychology of Cyber-Terrorism," 2003.
177. Rogers, "Psychology of Cyber-Terrorism," 2003.
178. Symantec, *Internet Security*, 2012.
179. Holt, "Attack Dynamics," 2011.
180. Rogers, "Psychology of Cyber-Terrorism," 2003.
181. Shaw, Ruby, and Post, "Insider Threat," 1998.
182. Shaw, "Role of Behavioral Research," 2006.
183. Shaw and Stock, "Behavioral Risk Indicators," 2011.
184. Shaw and Stock, "Behavioral Risk Indicators," 2011.
185. Shaw, Ruby, and Post, "Insider Threat," 1998.
186. Shaw, Ruby, and Post, "Insider Threat," 1998.
187. Shaw, Ruby, and Post, "Insider Threat," 1998.
188. Shaw, "Role of Behavioral Research," 2006.
189. Shaw and Stock, "Behavioral Risk Indicators," 2011.
190. Rogers, "Meaningful Hacker Taxonomy," 2005.
191. Andrew Bissett and Geraldine Shipton, "Some Human Dimensions of Computer Virus Creation and Infection," *International Journal of Human Computer Studies*, 52, No. 5 (2000), pp. 1071–5819.
192. Sarah Gordon, "The Generic Virus Writer," 4th International Virus Bulletin Conference, Jersey, U.K. (September 1994). http://vxheaven.org/lib/static/vdat/epgenvr2.htm
193. Sarah Gordon, "The Generic Virus Writer II," 6th International Virus Bulletin Conference, Brighton, U.K. (September 1996). www.research.ibm.com/antivirus/SciPapers/Gordon/GVWII.html
194. Sarah Gordon, "Virus Writers: The End of Innocence?" Presented at the 10th International Virus Bulletin Conference (September 2000). http://vxheaven.org/lib/asg12.htm
195. Bissett and Shipton, "Some Human Dimensions," 2000.
196. Bissett and Shipton, "Some Human Dimensions," 2000.

197. Bissett and Shipton, "Some Human Dimensions," 2000.
198. Gordon, "Generic Virus Writer," 1994.
199. Gordon, "Generic Virus Writer II," 1996.
200. Gordon, "Virus Writers," 2000.
201. Myers, *Social Psychology*, 2005.
202. Gordon, "Generic Virus Writer II," 1996.
203. Gordon, "Generic Virus Writer II," 1996.
204. Gordon, "Generic Virus Writer II," 1996.
205. Gordon, "Generic Virus Writer II," 1996.
206. Gordon, "Generic Virus Writer II," 1996.
207. Rogers, "Meaningful Hacker Taxonomy," 2005.
208. Gordon, "Virus Writers," 2000.
209. A. S. Dalal, & Raghav Sharma, "Peeping into a Hacker's Mind: Can Criminological Theories Explain Hacking?" *ICFAI Journal of Cyber Law* 6, no. 4 (2007): 34–47.
210. Rogers, "Meaningful Hacker Taxonomy," 2005.
211. Holt, Burruss and Bossler, "Social Learning and Cyber-Deviance," 2010.
212. Gordon, "Virus Writers," 2000.
213. Gordon, "Virus Writers," 2000.
214. Young, Zhang, and Prybutok, "Hacking into the Minds," 2007.
215. Shaw and Stock, "Behavioral Risk Indicators," 2011.
216. Shaw and Stock, "Behavioral Risk Indicators," 2011.
217. M. E. Kabay, "Time for Industry to Support Academic INFOSEC," M. E. Kabay's Website, 2004. www.mekabay.com/opinion/endowed_chairs.pdf

CHAPTER 13

THE INSIDER THREAT

Gary L. Tagg, CISSP

13.1 INTRODUCTION 13·1

13.2 THREATS FROM INSIDERS 13·2

13.2.1 How Common Are Insider Attacks? 13·2
13.2.2 Examples of Insider Attacks 13·3
13.2.3 Types of Insider Threats 13·5
13.2.4 Internet-Based Systems 13·6
13.2.5 Service Provider Threats 13·7
13.2.6 System Administration Threats 13·7

13.3 MITIGATING THE INSIDER THREAT 13·7

13.3.1 System and Asset Inventories 13·8
13.3.2 Data Loss Prevention (DLP) 13·8
13.3.3 Internal Honeypots 13·10

13.4 CONCLUDING REMARKS 13·10

13.5 FURTHER READING 13·11

13.6 NOTES 13·11

13.1 INTRODUCTION. An insider is someone who has been given a role within an organization and has access to premises and/or internal systems and information. There are many types of roles; some are core to the business and performed by employees, while others are non-core and contracted out to service providers such as cleaning, maintenance, or information technology (IT). The categories of insiders may be classified as:

Current staff work directly for and under the control of the organization's management. This category includes employees as well as temporary staff, contractors, and consultants. Most of these people are located on the organization's premises and are connected to the internal network with access to internal information.

Departing staff represent one of the highest risks to an organization. These people consist of employees who have resigned or are planning to do so, temporary staff, contractors or consultants coming to the end of their contract, as well as all the people whose employment or services are being ended by the organization. These people still have access to internal information, and may be motivated to take that information with them when they leave, or to commit sabotage in revenge for perceived wrongs.

Former staff are those who are no longer employed by or providing services to the organization. This group still has insider knowledge, and without mitigating controls, they can do substantial damage long after they have left the organization. Former staff can be highly motivated to attack their former employers.

Service-providers: Organizations have many roles that are necessary for the smooth running of the business but that are not core to the company's mission. Examples of these roles are cleaning services, maintenance, and IT. Over the last 20 years, service providers have steadily moved up the service stack to perform core functions as well. It is now common for service providers to perform business operations and even to be the first point of contact with customers. Although not all service providers will need access to premises or internal systems, there are many roles that do.

The key distinction between service providers and staff is that staff (who can be provided by service providers) are working under the control of the organization's management and subject to its policies and procedures; in contrast, service providers are providing a specified service, and all personnel are the responsibility of the service provider. This situation increases risk to the organization, as it has little control over service-provider staff, but the threats are the same as from internal staff.

Partners: Organizations are partners when they work together on a business venture. Access to information goes up to senior business leaders within the partners. Partnerships may be short or long term, and partners may be competitors at the same time as being partners. This situation creates risk for the organization, which has to share information relevant to the partnership but not other internal information.

The term *partner* is sometimes used where service providers are performing key functions for an organization and success of the service is essential. For these key contracts, it is good practice to treat the service provider more like a partner than a vendor, but for insider threat purposes they are service providers.

13.2 THREATS FROM INSIDERS. Successful organizations need to have people in many different roles, with the primary distinction between business roles (business management, sales, customer services, and product design) and support/infrastructure roles (IT, finance, logistics, human resources, and the like).

The people in each role need access to information and systems to perform their role, and the key point to make at this stage is that the impact an insider attack makes on the organization is related to the insider role. A salesman leaving an organization to join a competitor will have had access to customer- and product-related information, whereas someone in product design is likely to have access to valuable intellectual property on current and future products. Customer service center staff are likely to have access to customers' personally identifiable information (PII) that can be used to commit identity fraud.

IT is a high-risk area, and is essential to the efficient running of the business. Most of an organization's information is stored in its IT systems, and without proper controls, IT administrators could have access to everything.

13.2.1 How Common Are Insider Attacks? According to the 2011 CyberSecurity Watch Survey,[1] 21 percent of electronic crime events were conducted by insiders, 58 percent by outsiders, and 21 percent unknown. Of the insider events, 76 percent were dealt with internally without legal action or law enforcement, which is likely a reason that external attacks are more often in the news.

The 2012 U.K. Information Security Breaches Survey[2] provides some interesting statistics.

- 6 percent of the responding 447 organizations reported staff sabotage of systems, with some organizations experiencing incidents on a weekly basis

- 19 percent of large organizations reported that staff used systems to commit theft or fraud, and this figure had doubled since the survey in 2010
- With regard to other incidents caused by insiders, large organizations were more likely to experience incidents, with 82 percent reporting incidents versus 45 percent for small organizations
- 61 percent of large organizations reported unauthorized access to systems or data
- 45 percent reported a breach of data protection laws or regulations
- 36 percent reported misuse of confidential information
- 47 percent reported loss or leakage of confidential information

The 2013 U.K. Information Security Breaches Survey[3] includes the following relevant information for this chapter (quoting directly from the Executive Summary—note U.K. spelling):

- 36% of the worst security breaches in the year were caused by inadvertent human error (and a further 10% by deliberate misuse of systems by staff)
- 57% of small businesses suffered staff-related security breaches in the last year (up from 45% a year ago)
- 17% of small businesses know their staff broke data protection regulations in the last year (up from 11% [in the 2012 report])
- 14% of large organisations had a security or data breach in the last year relating to social networking sites
- 9% of large organisations had a security or data breach in the last year involving smartphones or tablets
- 4% of respondents had a security or data breach in the last year relating to one of their cloud computing services
- 4% of the worst security breaches were due to portable media bypassing defences

13.2.2 Examples of Insider Attacks. To get an understanding of the types of insider events, a number of organizations maintain databases and publish results. The FBI maintains an Insider Threat page on their Website[4] containing a list of prosecuted insider theft cases. A short summary of these cases follows:

- Retired research scientist conspired with current and former employees to steal trade secrets from his former employer, and sell them to companies in China.
- Employee working for two different U.S. companies stole trade secrets from both companies which were used to benefit Chinese universities.
- A research scientist stole trade secrets from her employer and made them available for sale through her own company.
- An employee stole customer and employee lists, contract information, and other trade secrets to provide to a foreign government, but instead gave them to an undercover FBI agent.
- A computer programmer working for a financial firm copied proprietary software during his last few days at the company.
- An employee who was fired had kept copies of trade secrets. These trade secrets were then sold to a rival company.

- An employee stole and attempted to sell trade secrets that provided everything needed to start a competing business.
- Over a two-day period, an employee copied hundreds of technical documents from her employer, which were found in her luggage during a check at the airport.
- Spies working for U.S. defense companies stole internal information about the space shuttle, Delta IV rocket, the C-17 military plane, and submarine propulsion systems, along with other information. This information was provided to the Chinese government.

In a summary of the "Top Five Insider Attacks of the Decade," the Linux.com editorial board listed the following cases[5]:

- Roger Duronio was convicted in 2006 to eight and a half years in federal prison[6] for his actions in 2002 when, apparently in a fit of pique at not receiving what he considered an adequate annual bonus, he sabotaged the computer systems of his employer, UBS PaineWebber. "UBS was hit on March 4, 2002, at 9:30 in the morning, just as the stock market opened for the day. Files were deleted from up to 2,000 servers in both the central data center in Weehawken, N.J., and in branch offices around the country. Company representatives never reported the cost of lost business but did say it cost the company more than $3.1 million to get the system back up and running. Duronio worked at UBS as a systems administrator until he quit a few weeks before the attack. Witnesses testified that he quit because he was angry that he didn't receive as large an annual bonus as he expected. Investigators found copies of the malicious code on two of his home computers and on a printout sitting on his bedroom dresser."[7]
- In 2005, a sting operation by a *The Sun* reporter from Britain netted him confidential details of more than a thousand "… accounts, passports, and credit cards …" from NatWest and Barclays banks. This led to investigations of call centers in India, when criminals "… boasted of being able to provide details of as many as 200,000 bank accounts in a month, which, he further said, came from more than one call center."[8]
- The "Athens Affair" was discovered when investigators looked into the apparent suicide of an electrical engineer, Costas Tsalikidis, in his Athens apartment. The inside job by Tsalikidis, the head of network planning, and unknown others compromised the Vodafone-Panafon company ("Vodafone Greece") using malware and may have resulted in monitoring and recording of conversations involving "… the prime minister, his defense and foreign affairs ministers, top military and law-enforcement officials, the Greek EU commissioner, activists, and journalists."[9]
- San Francisco network administrator Terry Childs locked other employees out of the city's network in July 2008 because he claimed that his supervisor was unqualified to have administrative control. He was sentenced to four years in California state prisons. "Prosecutors characterized the former network administrator as a power hungry control freak who couldn't be managed."[10]
- Bradley Manning, a 22-year-old U.S. Army intelligence analyst, was arrested in May 2010 and charged in July 2010 with leaking nearly half a million classified U.S. videos and cables to the WikiLeaks project.[11] He was charged with "aiding the enemy" in February 2012[12] and accused of aiding the terrorist group al

Qaida through his actions.[13] In January 2013, the trial was rescheduled until June 2013[14] and Manning was denied the opportunity to justify his actions using a whistleblower defense.[15]

- *The Sunday Times* reported in 2012 that "Confidential personal data on hundreds of thousands of Britons is being touted by corrupt Indian call centre workers, an undercover investigation has discovered. Credit card information, medical and financial records are being offered for sale to criminals and marketing firms for as little as 2p." Records of 500,000 Britons were apparently for sale, many of which were supposedly less than 72 hours old and included "… sensitive material about mortgages, loans, insurance, mobile phone contracts, and Sky Television subscriptions …"[16]

13.2.3 Types of Insider Threats. There are three main classifications of insider threats: accidental, malicious, and nonmalicious.

13.2.3.1 Accidental Threats. Accidental threats are generally caused by mistakes; for example, staff may not follow operating procedures due to carelessness, disregard for policies, or a lack of training and awareness of the right thing to do. An example is a customer service representative who accidentally breaches client confidentiality by emailing client information to the wrong email address. Such errors may be caused by the use of email clients that have an auto-complete feature on the email address; staff under pressure to keep up with the volume of work may not notice the error before they send the data. This error is a particularly high risk for financial services organizations where in some jurisdictions a client confidentiality breach is a criminal offense.

Other typical examples include a database administrator who accidentally deletes a database table during maintenance, a systems administrator allowing a programmer to modify a production system without proper approvals, or an operator reformatting a disk drive without having two full, verified backups.

13.2.3.2 Malicious Threats. Malicious threats deliberately try to damage the organization or to benefit the attacker. Disgruntled IT administrators can sabotage IT systems, bringing an organization to a halt. There have been many incidents where both current and former administrators have deliberately caused system issues for various motives: enjoying the lifestyle of traveling around the world in luxury to fix the problems they created, extorting money from the organization, or simply causing as much damage as possible.

Some company information is highly valuable and specifically targeted by attackers. PII is one category that is sometimes illegally copied by staff to conduct identity theft and fraud, or to sell it to criminals. Another category is intellectual property (IP) such as trade secrets. Staff may take this information to help them with their next job or to sell to competing companies. This crime is thought to be common with IT developers who often seek to take their source code with them. Industrial espionage sponsored by rival companies or foreign governments is another common threat to IP (see Chapter 11 in this *Handbook* for examples).

Information can leave an organization by being copied to removable storage such as USB flash or hard drives and CD/DVD writers. Portable 3TB drives are now readily available, which means entire databases can be copied to a drive measuring less than 7×5 inches in size. With gigabit ethernet becoming standard, the time required to copy

this data is rapidly decreasing. Other common channels include emailing attachments to external email addresses, uploading files to external email services and to Internet Websites, and using cloud backup and cloud storage tools. To address these data-leakage channels, products known as data loss prevention (DLP) systems are increasingly being installed in organizations.[17]

There are also the common physical threats, such as taking printed information from people's desks or from the office printers. Where logical access controls are strong, staff intent on stealing information can take photographs of documents or information on screen with the high-resolution cameras in today's mobile phones.

There are some incidents where the motive may be conscience based as well as having a desire to damage an organization. Since the financial crisis in 2008, governments have increased their efforts to reduce tax avoidance, and are aggressively pursuing the use of foreign tax havens. There have been a number of publicized incidents where insiders have provided lists of offshore clients and accounts to country tax authorities.[18,19]

13.2.3.3 Nonmalicious Threats. Nonmalicious threats are actions taken deliberately by people without intent to damage the organization. Often, the motive is to increase productivity, and the mistakes occur due to a lack of training or awareness of policies, procedures, and the risk. There have been many incidents in which staff loaded internal information onto Internet-based systems, some of which have no access controls. This error makes the information available to anyone who uses the sites and is often found and indexed by search engines.

One incident occurred in the early days of cloud computing. A drug researcher was given a lengthy lead time by his internal IT department for the delivery of infrastructure to conduct simulations. What he did instead was use his credit card to buy time on cloud-based systems and ran the simulations there instead. This would have put valuable intellectual property at risk had these Internet systems been compromised. This information was also sitting on the cloud provider's storage systems; what if the cloud vendor were to fail to reinitialize the storage when reallocating released storage to another customer, and the next user of the storage were to understand the value of what came pre-loaded on their system—and be dishonest?

Another example that is often reported in the media is the loss of PII when staff copy information to laptop computers or to removable storage devices such as USB drives or CDs/DVDs, which are then lost or stolen.[20]

One common insider threat that can happen for both malicious and nonmalicious reasons is to email internal information to their home email address. Once on the staff member's own computer, the information is vulnerable to theft, successful attack on the computer or email account, or recycling of the machine by donating it to charity, or giving it away to family or friends. There have been many media reports of people finding sensitive information on secondhand computer hard disks.[21]

The nonmalicious motive for this practice is often to enable the employee to work from home, or the information is needed for a business trip. The malicious motive is to take the information with them, for reasons covered earlier.

13.2.4 Internet-Based Systems. The growing trend to use outsourced applications available over the Internet rather than internal systems can contribute to insider crime. With internal systems, when someone leaves an organization, they no longer have physical access to premises, and their network and application user-ids are supposed to be disabled immediately, removing access to corporate networks, computers,

and applications. Appropriate action makes it difficult for even maliciously motivated former staff to access confidential systems and their information.

But with Internet-based systems, the former staff member may still have access to confidential data via any Internet connection. It is very difficult to ensure all application accounts are promptly closed when someone leaves, creating a high risk of continued access to these systems. The people concerned may be deterred by the risk of being prosecuted for misusing their rights after leaving the organization, but their risk can be reduced if they use a former colleague's log-on account.

Organizations don't have to avoid using these external services; indeed, these external services are often among the best available, and can be provided at a much lower cost than hosting internally. However, security officers must coordinate closely with their human resources departments and the IT people responsible for maintaining the lists of external providers to ensure that access by former employees to external services is shut down as quickly as access to internal systems.

13.2.5 Service Provider Threats. Even organizations with strong personnel controls that reduce the risk from internal staff may have no protection from external service-provider staff. Although organizations commonly include their policies, standards, and procedures into contracts and treat this precaution as sufficient to address risk, the service provider may not consistently perform all the required controls, either through attempts to maximize profit or through poor management.

For example, a service provider faced with high staff turnover might bypass preemployment checks to get replacement staff quickly onto a service, particularly if there are service performance issues. Similarly, poor standards for handling termination of employment at the service provider could lead to compromise of client information.

In cases where an entire IT service is being provided by a service provider, the service provider's staff and IT administrators may have access to all the client organization's information. Email is a particular risk that may contain a great deal of an organization's intellectual property and that has powerful search facilities to easily target individuals and specific information.

13.2.6 System Administration Threats. System administrators have privileged access to an organization's IT infrastructure and, in poorly managed systems, may have access to critical and sensitive information on the systems. System administrators may deliberately attack the availability of an organization's systems even if they cannot access the data themselves. Administrators may destroy individual data sets, applications, or entire systems and networks—and may be able to destroy system backups to make the organization's recovery more difficult. This risk is most acute in smaller organizations, in which one person may manage the servers, applications, networks, email, backups, and perhaps even user administration. In larger organizations, it is much easier to segregate these roles as they should be, thus limiting access and therefore the damage one person can cause.

Even in large organizations with thousands of servers with effective segregation implemented, an administrator may be able to submit a job that can be run simultaneously on every server to wipe every hard disk. This type of threat highlights the importance of controlling the access and scope of administrators and preventing unauthorized changes from being implemented.

13.3 MITIGATING THE INSIDER THREAT. This section describes at a high level a few of some specific mitigating controls and how they address the insider threat.

For more details on the wider range of controls, refer to the detailed information in other chapters in this *Handbook* such as Chapter 15 (Penetrating Computer Systems and Networks), Chapter 28 (Identification and Authentication), Chapter 31 (Content Filtering and Web Monitoring), Chapter 45 (Employment Practices and Policies), Chapter 52 (Application Controls), Chapter 53 (Monitoring and Control Systems), Chapter 54 (Security Audits and Inspections), Chapter 55 (Cyber Investigation), Chapter 56 (Computer Incident Response Teams), and Chapter 68 (Outsourcing and Security).

13.3.1 System and Asset Inventories. If you don't know what you have, you can't manage it. Without an inventory of servers, you can't ensure they are patched, enabling an insider to compromise them and use them for their own purposes. Without a list of applications running on each server, you may have unnecessary servers sitting on your network being used for unauthorized purposes. The active systems on the network also need to correlate to the inventory. An administrator who doesn't remove a redundant system in your Internet DMZ could use it to bypass all of your network perimeter controls after he has left the organization.

In particular, one of the most dangerous tools for insider crime is the unauthorized and undetected wireless access point—easily purchased from any electronic store for at low cost and capable of transferring data from an internal network to unauthorized devices within the corporate facilities or even to external agents within a modest radius outside the building. See Chapter 33 in this *Handbook* for discussion of wireless network security.

13.3.2 Data Loss Prevention (DLP). There have been technology adoptions over the last 10 years which have made it much easier for staff to either accidentally or deliberately breach the confidentiality of organizational and client information. Data loss prevention (DLP) is a class of IT security system increasingly being implemented by organizations to help address these risks. A typical DLP system needs to address the following vulnerabilities at a minimum:

- Mobile storage devices/removable media, such as USB memory sticks and hard drives, mobile phones, memory cards, CD/DVD writers, along with infrared, Bluetooth, FireWire, and SCSI-connected storage.
- File uploads—including encrypted data—to external Websites via the standard protocols within Web browsers, such as ftp, http, and https. These controls may be enforced at both the Internet gateway as well as on the desktop.
- Detection of when laptops are not on the corporate network and preventing files being copied to noncorporate file shares.

For the majority of staff, the DLP system can be configured to block attempts to copy and upload data to these data-leakage channels. However, for most of these channels, there are going to be some people who have a genuine business need to copy and upload information, resulting in exceptions to the policy. The DLP system can help manage this risk by creating a log of what has been copied to support incident investigation or to enable a review of what a staff member has copied out of the organization. These facilities can be particularly useful when an employee hands in his resignation.

13.3.2.1 Email Data Leakage. Once the low-hanging fruit of removable storage and file uploads are blocked, staff begin to export information via email. The motive may not be malicious, but it still results in the organization's losing control over the information. To address this risk, the DLP system can be configured to report on people sending attachments to home email addresses, which is the usual destination for information, or to any unauthorized email address. These reports can then be used to increase staff awareness about policy or to support disciplinary processes for deliberate data leakage.

However, the true benefit from DLP comes when the business areas are engaged, because DLP can drive a process of identifying and defining the critical information that has to be protected. For example, the IT development group could have a DLP rule that blocks emails containing source code, to prevent developers taking their work with them when they leave. Even if source code is placed in an encrypted zip file to try and avoid detection, the metadata (e.g., filenames and identifiers of file creators) may still be readable and can trigger the rule.

For other areas of the business, client information, business strategy, business results, intellectual property, and PII such as credit card numbers and Social Security numbers can be configured into the DLP system.

13.3.2.2 Data Leakage Using Cloud Storage. Widespread availability of external data storage facilities (e.g., Dropbox and Google Drive) adds to the complexity of DLP. Careful application of Web monitoring and blacklisting specific URLs may be helpful, but determined opponents of the regulations may circumvent such methods using a variety of proxy avoidance Websites which mask the destination of the HTTP request. Security administrators should be on the lookout for new sites so they can add them to the corporate blacklists for outbound communications through their firewalls.

13.3.2.3 Difficulties with DLP. DLP is not a panacea. It is difficult if not impossible to prevent someone determined to leak data, but with a DLP system, one can make it difficult for them to do so without detection, and this barrier usually deters the majority of people.

With email DLP, it is very difficult to identify safe blocking rules that prevent data leakage. For example, if employees routinely send emails to customers, some of these customers are going to be using their home email addresses, which means that a hard block (blacklisting) of emails sent to home email addresses won't be feasible.

13.3.2.4 Legal Issues with DLP. With the implementation of DLP, an organization progresses from investigating reported incidents to actively monitoring for breaches of company policy. One of the major issues with implementing a global DLP system is that active monitoring may be subject to privacy and workplace laws, and failure to comply with these laws in some countries is a criminal offence.

Additionally, with customer and organization information being captured in DLP logs, consideration needs to be given as to where the log files are stored and who will be reviewing them. As an example from the financial services industry, some information is price sensitive and there are strict requirements on who can access it, to prevent insider dealing. Client information within the log files may be covered under banking secrecy laws and regulations. If the log files for one country are stored on a server in another country, then outsourcing regulations need to be complied with as well, and planners must identify the most stringent regulations to ensure efforts for

compliance. Data protection laws have been enacted in many countries that require data subjects to give their agreement to their information being processed and for defined purposes; therefore, monitoring of communications may need to be included in customer contracts and other data protection declarations.

Despite all of these hurdles, for most countries, it is possible to roll out a DLP system. As part of the project planning stage, the organization needs to commission a legal investigation for each country to understand whether it is lawful to actively monitor corporate email, but more importantly, the preconditions for any monitoring to lawfully take place. For countries that forbid email monitoring or any form of workplace surveillance, it is usually possible to implement blocking controls on writing to removable storage along with uploads to external Websites, provided no log files are kept.

13.3.2.5 Remote Access Solution. A commonly implemented remote access solution is the provision of Webmail from home computers. With Webmail you need to configure the system to prevent staff opening attachments within local applications on the home computer. If this is not blocked, staff can copy corporate information by saving the open attachment to their local hard drive. In addition, any use of the Remote Desktop Protocol (RDP) also needs to be properly configured to prevent local USB resources being mapped to the remote computer.

13.3.3 Internal Honeypots. Honeypots are attractive systems or nodes that appear to have valuable confidential data. Roger Grimes, author of a textbook about honeypots,[22] writes:

> One of the best things you can do to get early warning of internal attackers is to implement honeypots. … because they're low cost and low noise—and they work!
>
> Because the internal attackers you seek could be trusted IT employees, the entire project must be kept secret. It should be known only to the sponsor, management, and the implementers. Often, we give the project a boring code name like Marketing Business Development, which is used in all documents and e-mails, avoiding terms having anything to do with honeypots. Don't even tell the network security people about it, in so far as you can and still have an operational project. Then take a few computers that are destined for de-provisioning or the scrap heap and turn them into your honeypots.[23]

The point of such a honeypot is that there should be zero access to it: It serves no function and there is never a reference to it in any internal or external documentation. Thus, anyone who accesses it is either doing, so by pure accident or is violating policies governing unauthorized access to internal data (there *is* no one authorized to access the honeypot). Detailed logging should be in place at all times to provide forensic information immediately; however, administrators should ensure that they can actually visualize exactly what is being done in real time, not simply rely on the detailed log files for after-the-fact analysis. An early-warning system should immediately alert system administrators to the problem (preferably as the access is in progress) via screen messages, voice messages, and text messages.

13.4 CONCLUDING REMARKS. Fighting the insider threat need not be solely reactive. In addition to the entire spectrum of information assurance measures covered throughout this *Handbook*, maintaining an environment of security awareness and of encouragement, trust, fair-dealing, and long-term commitment to the welfare of employees must remain among the very best approaches to reducing the risk of insider crime.

13.5 FURTHER READING

Department of Defense. *DoD Insider Threat Mitigation: Final Report of the Insider Threat Integrated Process Team.* DoD. 2000. www.dtic.mil/cgi-bin/GetTRDoc?AD = ADA39138

Noonan, T. & E. Archuleta. *The National Infrastructure Advisory Council's Final Report and Recommendations on the Insider Threat to Critical Infrastructures.* NIAC. 2008. www.dhs.gov/xlibrary/assets/niac/niac_insider_threat_to_critical_infrastructures_study.pdf

Silowash, G., D. Gappelli, A. Moore, R. Trzeciak, T. J. Shimeall, L. Flynn. *Common Sense Guide to Mitigating Insider Threats,* 4th ed. CMU/SEI. 2012. Technical Report CMU/SEI-20120TR-012. www.sei.cmu.edu/reports/12tr012.pdf

13.6 NOTES

1. CMU/SEI 2011
2. PwC 2012
3. PwC 2013
4. See www.fbi.gov/about-us/investigate/counterintelligence/the-insider-threat
5. Linux.com editorial staff, "Top Five Insider Attacks of the Decade," *LINUX.COM*, January 13, 2011, www.linux.com/news/technology-feature/security/397143-top-five-insider-attacks-of-the-decade
6. Sharon Gaudin, "Ex-UBS Systems Admin Sentenced To 97 Months In Jail," *InformationWeek*, December 13, 2006, www.informationweek.com/ex-ubs-systems-admin-sentenced-to-97-mon/196603888
7. Sharon Gaudin, "Ex-UBS Sys Admin Found Guilty, Prosecutors To Seek Maximum Sentence," *InformationWeek*, July 19, 2006, www.informationweek.com/ex-ubs-sys-admin-found-guilty-prosecutor/190700064
8. SiliconIndia News, "Indian call centers selling U.K.'s secrets," *SiliconIndia News*, June 23, 2005, www.siliconindia.com/shownews/Indian_call_centers_selling_UKs_secrets-nid-28560-cid-2.html
9. Vassilis Prevelakis and Diomidis Spinellis, "The Athens Affair: How some extremely smart hackers pulled off the most audacious cell-network break-in ever," *IEEE SPECTRUM*, June 29, 2007, http://spectrum.ieee.org/telecom/security/the-athens-affair
10. Robert McMillan, "Network admin Terry Childs gets 4-year sentence," *NetworkWorld*, August 17, 2012, www.networkworld.com/news/2010/080710-network-admin-terry-childs-gets.html
11. Chris McGreal, "US private Bradley Manning charged with leaking Iraq killings video," *The Guardian*, July 6, 2010, www.guardian.co.uk/world/2010/jul/06/bradley-manning-charged-iraq-killings-video
12. Karen McVeigh, "Bradley Manning defers plea after being formally charged with aiding the enemy: No date set for WikiLeaks suspect's trial but Manning's lawyer says he would object to any delay in the trial beyond June," *The Guardian*, February 23, 2012, www.guardian.co.uk/world/2012/feb/23/bradley-manning-defer-plea-charges
13. Associated Press, "Bradley Manning aided al-Qaida with WikiLeaks documents, military says: Manning, charged with aiding the enemy, accused of indirectly

aiding terrorist group by leaking thousands of documents," *The Guardian*, March 3 2012, www.guardian.co.uk/world/2012/mar/15/bradley-manning-wikileaks

14. Adam Gabbatt and Ed Pilkington, "Bradley Manning trial delayed until June after sentence reduction granted: Judge reschedules US soldier's trial to give more time for review of classified information related to WikiLeaks case," *The Guardian*, January 9, 2013, www.guardian.co.uk/world/2013/jan/09/bradley-manning-trial-delayed
15. Ed Pilkington, "Bradley Manning denied chance to make whistleblower defence: Judge rules that Manning will not be allowed to present evidence about his motives for the leak—a key plank of his defence," *The Guardian*, January 17, 2013. www.guardian.co.uk/world/2013/jan/17/bradley-manning-denied-chance-whistleblower-defence
16. Tom Gardner, "Indian call centres selling YOUR credit card details and medical records for just 2p," *MailOnline*, March 18, 2012, www.dailymail.co.uk/news/article-2116649/Indian-centres-selling-YOUR-credit-card-details-medical-records-just-2p.html
17. Eric Ouellet, "Magic Quadrant for Content-Aware Data Loss Prevention," Gartner, Inc., January 3, 2013, https://www.ca.com/us/register/forms/collateral/Magic-Quadrant-for-Content-Aware-Data-Loss-Prevention-2013.aspx
18. BBC 2012
19. BBC 2011
20. BBC 2008
21. BBC 2009
22. Roger A. Grimes, *Honeypots for Windows*, Apress, 2005.
23. R. A. Grimes, "Honeypots: A sweet solution to the insider threat: A honeypot can be a cheap, easy, and effective warning system against the trusted insider gone bad," *InfoWorld*, May 1, 2009, www.infoworld.com/d/security-central/honeypots-sweet-solution-insider-threat-922?page=0,0

CHAPTER 14

INFORMATION WARFARE

Seymour Bosworth

14.1 INTRODUCTION 14·2

14.2 VULNERABILITIES 14·2

14.2.1 Critical Infrastructure 14·2
14.2.2 Off-the-Shelf Software 14·3
14.2.3 Dissenting Views 14·4
14.2.4 Rebuttal 14·4

14.3 GOALS AND OBJECTIVES 14·5

14.3.1 Infrastructure 14·5
14.3.2 Military 14·5
14.3.3 Military Offensives 14·8
14.3.4 Government 14·8
14.3.5 Energy Systems 14·9
14.3.6 Transportation 14·11
14.3.7 Commerce 14·12
14.3.8 Financial Disruptions 14·13
14.3.9 Medical Security 14·14
14.3.10 Law Enforcement 14·15
14.3.11 International and Corporate Espionage 14·15
14.3.12 Communications 14·16
14.3.13 Destabilization of Economic Infrastructure 14·16

14.4 SOURCES OF THREATS AND ATTACKS 14·17

14.4.1 Nation-States 14·17
14.4.2 Cyberterrorists 14·19
14.4.3 Corporations 14·21

14.5 WEAPONS OF CYBERWAR 14·21

14.5.1 Denial of Service and Distributed Denial of Service 14·21
14.5.2 Malicious Code 14·22
14.5.3 Cryptography 14·23
14.5.4 Psychological Operations 14·23
14.5.5 Physical Attacks 14·24
14.5.6 Biological and Chemical Weapons and Weapons of Mass Destruction 14·25
14.5.7 Weapons Inadvertently Provided 14·25

14.6 DEFENSES 14·25

14.6.1 Legal Defenses 14·25
14.6.2 Forceful Defenses 14·26
14.6.3 Technical Defenses 14·27
14.6.4 In-Kind Counterattacks 14·27
14.6.5 Integration of Cyberwarfare into Military Planning 14·27
14.6.6 Cooperative Efforts 14·28

14.7 SUMMARY 14·29

14.8 FURTHER READING 14·29

14.9 NOTES 14·30

Information warfare is the offensive and defensive use of information and information systems to deny, exploit, corrupt, or destroy, an adversary's information, information-based processes, information systems, and computer-based networks while protecting one's own. Such actions are designed to achieve advantages over military or business adversaries.

—Dr. Ivan Goldberg, Institute for Advanced Study of Information Warfare

14.1 INTRODUCTION. Until recently, warfare was conducted by armed forces representing adversarial nations, or by revolutionary elements opposing their own governments. Today, although such conflicts still exist around the world, the ubiquitous nature of computers and associated technology has created new forces, new threats, new targets, and an accompanying need for new offensive and defensive weapons. Information warfare (IW), also known as *e-warfare* or *cyberwar*, is actually, or potentially, waged by all U.S. armed forces and by those of other nations, as well as by commercial enterprises, by activist groups, and even by individuals acting alone.

Conventional wars, whether large or small, are regularly reported by the news media. Information war, however, are largely ignored except by those with a professional interest in the field. One reason for this is that conventional warfare is a matter of life or death; photos and eyewitness accounts are dramatic reminders of human cruelty and mortality. In contrast, IW has so far been conducted bloodlessly, with only economic and political consequences. However, it is becoming increasingly evident that IW may soon be conducted in ways that could equal or exceed the death and destruction associated with conventional weapons.

Conventional wars are fought by known combatants with clearly defined allies and antagonists, but IW often is waged by unknown entities with uncertain allegiances and goals. IW may be conducted on many fronts simultaneously, with wars fought within wars, and with both civilian and military targets devastated.

The motives for conventional warfare were almost always territorial, religious, political, or economic. These are still important, but to them must be added the psychological motivations of groups and individuals—groups far more widely distributed and less easily overcome.

This chapter discusses information warfare in terms of the vulnerabilities of targets, participants' objectives, sources of threats and attacks, weapons used, and defenses against those weapons.

14.2 VULNERABILITIES. Until recently, concerns over the security of the technological infrastructure in technologically advanced nations have been viewed with skepticism. However, by the mid-1990s, opinion leaders in government, industry, and the security field were coming to grips with widespread vulnerabilities in the critical infrastructure.

14.2.1 Critical Infrastructure. In 1998, President Bill Clinton circulated Presidential Decision Directive 63, which outlined his administration's policy on critical infrastructure protection:

> Critical infrastructures are those physical and cyber-based systems essential to the minimum operations of the economy and the government. … They include, but are not limited to, telecommunications, energy, banking and finance, transportation, water systems and emergency services, both government and private.[1]

Having defined the very broad, vital areas that require protection, the paper went on to describe succinctly their vulnerability:

> The United States possesses both the world's strongest military and its largest national economy. Those two aspects of our power are mutually reinforcing and dependent. They are also increasingly reliant upon certain critical infrastructures and upon cyber-based information systems. …
>
> Because of our military strength, future enemies, whether nations, groups or individuals, may seek to harm us in non-traditional ways including attacks within the United States. Our economy is increasingly reliant upon interdependent and cyber-supported infrastructures and non-traditional attacks on our infrastructure and information systems may be capable of significantly harming both our military power and our economy.

A few examples of specific weaknesses were given by Jack L. Brock, Jr., director, Government-wide and Defense Information Systems, United States General Accounting Office:

> In May 1999 we reported that, as part of our tests of the National Aeronautics and Space Administration's (NASA) computer-based controls, we successfully penetrated several mission-critical systems. Having obtained access, we could have disrupted NASA's ongoing command and control operations and stolen, modified, or destroyed systems software and data.
>
> In August 1999, we reported that serious weaknesses in Department of Defense (DOD) information security continue to provide both hackers and hundreds of thousands of authorized users the opportunity to modify, steal, inappropriately disclose, and destroy sensitive DOD data.[2]

Although these "attacks" were carried out one at a time, and without malicious intent, it is apparent that they, and many others, could have been launched simultaneously and with intent to inflict the maximum possible damage to the most sensitive elements of the national infrastructure.

In a memorandum to its chairman, describing a report of the Defense Science Board Task Force on Defensive Information Operations, Larry Wright stated in 1999 that:

> The threats to the DoD infrastructure are very real, non-traditional and highly diversified. … The vulnerabilities of these United States are greater than ever before, and we know that over twenty countries already have or are developing computer attack capabilities. Moreover, the Department of defense should consider existing viruses and "hacker" attacks to be real "Information Operations or Warfare," what early aviation was to Air Power. In other words, we have not seen anything yet![3]

The report concluded that "[i]t is the view of this task force that DoD cannot today defend itself from an Information Operations attack by a sophisticated nation state adversary."

14.2.2 Off-the-Shelf Software. One characteristic of almost all military and civilian infrastructures is that they share, with more than 100 million computers, a single ubiquitous operating system, and many of the same applications programs, such as word processors, spreadsheets, and database software. These commercial off-the-shelf (COTS) products are available around the world, to friend and foe alike, and they appear to be more intensively studied by malefactors than by their security-inadequate producers. Each of these products presents entry points at which one common vulnerability may be exploited to damage or destroy huge portions of the national infrastructure.

Until, and unless, this software is rendered significantly more resistant to attack, all of its users remain at risk.

14.2.3 Dissenting Views. Not every influential observer concurs in these possible scenarios. Dr. Thomas P. M. Barnett, a professor and senior decision researcher at the Decision Support Department, Center for Naval Warfare Studies, U.S. Naval War College, voiced a fairly typical disagreement in 1999:

> If absence makes the heart grow fonder, network-centric warfare is in for a lot of heartbreak, because I doubt we will ever encounter an enemy to match its grand assumptions regarding a revolution in military affairs. The United States currently spends more on its information technology than all but a couple of great powers spend on their entire militaries. In a world where rogue nations typically spend around $5 billion a year on defense, NCW is a path down which only the U.S. military can tread.[4]

14.2.4 Rebuttal. It may be of some benefit to have spokespersons for this unworried viewpoint, but their opinions must be weighed against those, for example, of Scott Henderson, of the Navy-Marine Corps intranet, who said: "One of our critical capabilities will be how we are to defend our information and our information systems from an adversary's attack."[5] He stated that successful intrusions, or attacks, on Navy computer systems increased from 89 in 2000 to 125 by mid-2001, an annualized increase of 80 percent. Those figures did not include successful attacks that went undetected or unsuccessful attempts that may have identified a weak point from which to launch future and probably more successful, attacks.

A highly significant factor in IW is its asymmetric nature. The barriers to entry for attackers are low; their weapons can be inexpensive, easily obtained, highly effective, easily automated, and used with negligible risk of personal harm. In contrast, defensive measures are extremely costly in time, money, and personnel and they may be ineffective against even unsophisticated attackers using obsolete computers.

The 2013 *Annual Report to Congress: Military and Security Developments Involving the People's Republic of China* from the U.S. Office of the Secretary of Defense includes the following evaluation of China's "Anti-Access/Area Denial (A2/AD)" capabilities:

> An essential element, if not a fundamental prerequisite, of China's emerging A2/AD regime is the ability to control and dominate the information spectrum in all dimensions of the modern battlespace. PLA authors often cite the need in modern warfare to control information, sometimes termed "information blockade" or "information dominance," and to seize the initiative and gain an information advantage in the early phases of a campaign to achieve air and sea superiority. China is improving information and operational security to protect its own information structures, and is also developing electronic and information warfare capabilities, including denial and deception, to defeat those of its adversaries. China's "information blockade" likely envisions employment of military and non-military instruments of state power across the battlespace, including in cyberspace and outer space. China's investments in advanced electronic warfare systems, counter-space weapons, and computer network operations (CNO)—combined with more traditional forms of control historically associated with the PLA and CCP systems, such as propaganda and denial through opacity, reflect the emphasis and priority China's leaders place on building capability for information advantage.[6]

Considering the nature and extent of already successful attacks against major elements of U.S. military and civilian infrastructures, there appears to be no justification for discounting the views of those who believe that IW, in both its offensive and defensive roles, must be accorded the attention that surrounds any potentially cataclysmic force. This *Handbook*, especially Chapters 16, 17, 18, 20, and 21, contains many

examples of viruses, worms, and other malware that have created massive disruptions in very large networks. The worst-case scenarios presented here should serve to awaken a measured response in those who may have been unaware or unconcerned.

14.3 GOALS AND OBJECTIVES. Attacking forces, in information warfare, will always have a variety of strategic and tactical motives behind their actions; defensive forces generally have only one tactical goal—to blunt the enemy's attack and, if possible, to counterattack. Only after this is accomplished, and the nature of the attackers has been studied, can strategies for long-range operations be determined and effected.

14.3.1 Infrastructure. Depending on the target, an attacker's goals may vary widely, but attackers generally want to damage, subvert, or destroy the infrastructure. In doing so, an attacker would hope to bring government, the economy, and military operations to a standstill or at least to reduce their efficiency and effectiveness to instill fear, uncertainty, and doubt (FUD), and ultimately to induce widespread chaos that could cost many lives.

Although this view is entirely appropriate to wars between nations or to campaigns by terrorists, it must be tempered when considering commercial warfare, whose main goal is competitive financial advantage.

14.3.2 Military. Today, information warfare is a vital concern of area commanders under battlefield conditions. They must obtain complete, accurate, and timely information about their opponents' actions, intentions, weaknesses, and resources while denying the same to their adversaries. The ultimate objective for all of these activities is to support the military tactics that will maximize the enemy's body count, or at least to render its defenses ineffective, so that surrender becomes the only viable option. The other side of the coin, defensive tactics, are aimed at preventing enemies from accomplishing their objectives.

In the United States, the Joint Chiefs of Staff (for Army, Navy, Marine Corps, Coast Guard, and Air Force) have formulated the *Joint Doctrine for Operations Security* to be followed by all commanders of combatant commands in planning, preparation, and execution of joint operations. The publication states:

> Operations Security (OPSEC) is a process of identifying critical information and subsequently analyzing friendly actions attendant to military operations and other activities, to: (a) identify those operations that can be observed by adversary intelligence systems; (b) determine what indicators adversary intelligence systems might obtain that could be interpreted or pieced together to derive critical information in time to be useful to adversaries; and (c) select and execute measures that eliminate or reduce to an acceptable level the vulnerabilities of friendly actions to adversary exploitation.[7]

OPSEC is a process that could be applied to every element of civilian infrastructure, as well as to the military, although all sources of information commonly used by the military are not available to the civilian sector. Other military code words for intelligence activities are:

- HUMINT (human intelligence) is the most widely used source of information, as it has always been for both the civilian and military sectors. HUMINT is often the only source capable of direct access to an opponent's plans and intentions. Some

intelligence gathering is quite open, but covert or clandestine operations must be conducted in secrecy, so as to protect the sources of confidential information.

- SIGINT (signals intelligence) is obtained from communications (COMINT), electronics (ELINT), and foreign instrumentation signals (FISINT).
- COMINT (communications intelligence) is information intended for others and intercepted without leaving a trace.
- ELINT (electronic intelligence) derives technical or geographic location data from an opponent's electromagnetic radiations, other than those that arise from communications or from nuclear detonations or radioactive sources. The primary ELINT sources are radars (radio detection and ranging).
- FISINT (foreign instrumentation signals intelligence) is obtained from intercepting and analyzing metered performance parameters electronically transmitted from sources such as a ballistic missile.
- MASINT (measurement and signatures intelligence) is scientific and technical in nature. Its purpose is to identify distinctive features associated with a source, emitter, or sender so as to facilitate subsequent identification or measurement. These features include wavelength, modulation, time dependencies, and other unique characteristics derived from technical sensors.
- IMINT (imagery intelligence) is produced by photography, infrared sensors, lasers, radars, and electro-optical equipment. This equipment, operated from land, sea, air, or space platforms, provides strategic, tactical, and operational information.
- TECHINT (technical intelligence) is derived from the exploitation and analysis of captured or otherwise acquired foreign equipment.
- OSINT (open source intelligence) is available to the general public from news media, unclassified government publications, public hearings, contracts, journals, seminars, and conferences. The World Wide Web has become an important tool of OSINT.

The *Joint Doctrine for Operations Security* lists several generic military activities with some of their associated critical information. It must be the objective of all information warfare to acquire this critical information about their opponents while denying such information to them:

- *Diplomatic negotiations* include military capabilities, intelligence verification, and minimum negotiating positions.
- *Political-military crisis management* includes target selection, timing considerations, and logistic capabilities and limitations.
- *Military intervention* requires information about intentions, military capabilities, forces assigned and in reserve, targets, and logistic capabilities and constraints.
- *Counterterrorism* involves forces, targets, timing, strategic locations, tactics, and ingress and egress methods.
- *Open hostilities* information involves force composition and disposition, attrition and reinforcement, targets, timing, logistic constraints, and location of command and control (C^2) nodes.

- *Mobilization* requires information about an intent to mobilize before public announcement, impact on military industrial base, impact on civilian economy, and transportation capabilities and limitations.
- *Intelligence, reconnaissance, and surveillance* information includes purpose and targets of collection, timing, capabilities of collection assets, and processing capabilities.

In addition to the Joint Chiefs' doctrines, the Department of Defense and each individual branch of service have been charged with the responsibility for establishing task forces, advisory groups, training and awareness programs, and virtual information networks to mobilize IW forces and to bring into being a strong defense against enemy attack.

Further evidence of the importance of military information and the vulnerabilities that exist at this time is contained in the 2001 report of the Secretary of Defense to the President and the Congress:

> Information superiority is all about getting the right information to the right people at the right time in the right format while denying adversaries the same advantages. The United States enjoys a competitive advantage in many of the technical components of information superiority, but the U.S. also has vulnerabilities stemming from its increasing dependence on high technology. Experiences from Somalia to the Balkans have shown that low technology adversaries also can wage effective information campaigns, especially in urban environments.
>
> In the Information Age, the opportunities and obstacles to achieving national security objectives often are informational in nature. Information superiority is a principal component of the transformation of the Department. The results of research, analyses, and experiments, reinforced by experiences in Kosovo, demonstrate that the availability of information and the ability to share it significantly enhances mission effectiveness and improves efficiencies. Benefits include: increased speed of command, a higher tempo of operations, greater lethality, less fratricide and collateral damage, increased survivability, streamlined combat support, and more effective force synchronization. Kosovo also highlighted the shortage of assets for intelligence, surveillance, and reconnaissance, as well as the need for more secure interoperability and information protection, especially within coalitions.
>
> To ensure that the above prerequisites are in place, DoD is developing appropriate policy and oversight initiatives, actively pursuing opportunities to improve international cooperation in the areas of Command, Control, Communication, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) and space-related activities, partnering with industry, and working to anticipate and understand the implications of emerging information technologies.
>
> The quality of DoD's infostructure will be a pacing item on the journey to the future. The ability to conceive of, experiment with, and implement new ways of doing business to harness the power of Information Age concepts and technologies depends upon what information can be collected, how it can be processed, and the extent to which it can be distributed. The ability to bring this capability to war will depend upon how well it can be secured and its reliability. DoD envisions an infostructure that is seamless with security built-in, one that can support the need for increased combined, joint, and coalition interoperability, leverages commercial technology, and accommodates evolution.[87]

Although not as well publicized as are the U.S. defensive efforts, equal attention, time, and resources are being expended on actual and possible offensive operations. Every objective, every tactic, and every recommendation just mentioned, and some too sensitive to discuss here, are subjects for study and implementation of offensive strategies and tactics aimed at enemies, present and future.

14.3.3 Military Offensives. The nature of U.S. offensive planning has been described in detail in a top-secret presidential memo published by *The Guardian* (U.K.) on June 7, 2013.

The memo states that:

> The United States has an abiding interest in developing and maintaining use of cyberspace as an integral part of U.S. national capabilities to collect intelligence, and to deter, deny, or defeat any adversary that seeks to harm U.S. national interests in peace, crisis, or war. … The United States Government shall conduct DCEO [Defensive Cyber Effects Operations] and OCEO [Offensive Cyber Effects Operations] under this directive consistent with its obligations under international law, including with regard to matters of sovereignty and neutrality, and as applicable, the laws of armed conflict. This directive pertains to cyber operations, including those that support or enable kinetic, information, or other types of operation.[9]

In this context, "kinetic operations" are a recognized euphemism for warfare, and this memo provides the U.S. Government with a broad mandate to protect U.S. interests, as it sees them, with any means available to it without a declaration of war.

At the time of this writing (July 2013), the full implications of this serious security breach cannot be evaluated. At the very least, it will provide the critics of U.S. policy with reinforcement for their view of the United States as a militaristic, even terrorist force in international relations. For defenders of U.S. policies, it will represent an ordered, rational, response to threats that must be contemplated and guarded against.

14.3.4 Government. The objectives of government, at every level, must be to protect the lives and welfare of its constituencies. Any breakdown in an essential government function may produce marked unrest, rioting, vandalism, civil disobedience, and possibly much bloodshed.

Just as in the military, government must be able to defend itself against an information attack waged by any enemy of the established order. Although not every element of government is perceived by all to perform a useful function, there are agencies without which it would be virtually impossible to sustain a developed nation's day-to-day activities.

At the federal level, civil servants' salaries, Social Security payments, tax collections and disbursements, military expenditures, lawmaking, and a myriad of other functions and activities can be carried out only with the active and pervasive use of computers and computer networks. In the past, some of these computer operations have been penetrated by hackers, crackers, and political dissidents, but only one at a time. It does not require a science fiction writer to imagine what the effect would be if simultaneous attacks were successfully launched against major federal government agencies.

At state levels, although the effects would be more constrained geographically, a great deal of damage could be done to emergency response units, to police and judiciary functions, and to health and welfare services. All of these depend on computerized functions that are protected even less than those of federal agencies.

For municipalities and even smaller governments, zoning enforcements and other local functions can be suspended without serious consequences, but police radio and computer networks are easily penetrated, and their ability to maintain law and order compromised.

As demonstrated by many previous incidents, government functions at any level are susceptible to information warfare. Natural events, Murphy's law (what *can* go wrong *will* go wrong), poorly configured systems, flawed operating systems and application programs, together with inadequate security measures underlie the vulnerability of government systems.

14.3.5 Energy Systems.[10] Oil, coal, and hydroelectric systems are critical elements of the national infrastructure. Supervisory control and data acquisition (SCADA) security in the electric power industry suffers from widespread misconceptions and a breakdown in communications between administrators and security experts. In brief,

- Attacks on electric power plants and the distribution grid may not result in the catastrophic scenarios painted by the promoters of panic, but any interruption in electric power delivery can cause widespread infrastructure disruption.
- SCADA systems controlling electric generators and distribution systems are not, in fact, isolated by air gaps from the Internet.
- On the contrary, vulnerability analysis teams have systematically and repeatedly demonstrated that power companies are unaware of the reality of their interconnectedness and vulnerabilities.
- There are documented cases of industrial espionage, sabotage, denial of service, and malware attacks on electric power grid SCADA systems.
- SCADA systems have been considered too stable to bother updating with current patches; as a result, they are consistently vulnerable to exploits of current (and even ancient) vulnerabilities.
- Many SCADA systems were developed without consideration of security, secure coding, or integration of security dimensions of software quality assurance.
- Government and academia have significant projects in place to advance SCADA security, but acceptance by industry is modest at best. Academics engaged in SCADA security research are doing a good job of reaching other academics through peer-reviewed presentations at academic conferences; they are less successful in reaching managers at power companies.
- Pressure is rising in the public sphere, in government circles, among security practitioners, and within the electric power industry to come to grips with the need for improved cybersecurity.
- The electric power industry must coordinate its efforts to implement well-established standards for protecting computer systems and networks in all its SCADA systems and related networks. In addition, the industry should implement cyber situational awareness solutions to integrate multiple inputs from SCADA and network sensors that will permit intelligent, agile response to attacks and effective forensic analysis of those attacks.

The electric power industry has become a fundamental underpinning of twenty-first century life. In a landmark report on "The Electricity Economy," author Jesse Berst and colleagues describe the convergence of growing demand, an increasing dependence on computerized SCADA systems, and the inevitable complexity of interactions among elements controlled by diverse entities with limited coordination.[11] To illustrate the

growth in electricity demands, the report's Table 1 shows global electricity demands of 2.06 terawatts (TW) in 1950 versus 3.8 TW in 2000 and a predicted 6.99 TW in 2050. The proportion of electricity as a percentage of global energy utilization was 10.4 percent in 1950 and 25.3 percent in 2000; by 2050 it may reach 33.7 percent. The authors add,

> Today we depend on electricity for basic needs such as food, water, shelter, communication, employment, and health care. Those needs are served by infrastructures for food preservation, water treatment, heat and light, phone service, Internet, offices, factories, hospitals and emergency response, to name a few. Yet all of those essentials degrade or disappear without electricity.[12]

In October 1997, the President's Commission on Critical Infrastructure Protection (the "Marsh Report" named after Commission Chairman Robert T. Marsh) included the following warning:

> Prolonged disruption in the flow of energy would seriously affect every infrastructure.
>
> The significant physical vulnerabilities for electric power are related to substations, generation facilities, and transmission lines. Large oil refineries are also attractive targets. The increase in transportation of oil via pipelines over the last decade provides a huge, attractive, and largely unprotected target array. Oil and gas vulnerabilities include lines at river crossings; interconnects; valves, pumps, and compressors; and natural gas city gates. Large metropolitan areas could be deprived of critical fuel for an extended period by a properly executed attack.
>
> The widespread and increasing use of Supervisory Control and Data Acquisition (SCADA) systems for control of energy systems provides increasing ability to cause serious damage and disruption by cyber means. The exponential growth of information system networks that interconnect the business, administrative, and operational systems contributes to system vulnerability.[13]

Electrical power systems have been harmed through data leakage, industrial espionage, insider threats and sabotage.[14] Incidents described in the reference include

- 2006 Japan's Power Plant Security Info Leaked Onto Internet
- 2007 Egypt Accuses Nuclear Employee of Spying
- 2007 Former Nuclear Plant Engineer Allegedly Took Data to Iran
- 2007 Saboteur of California Power Grid Gained Access Despite Warning
- 2009 Fired Nuclear-Power-Plant Employee Arrested for Hacking Systems
- 2009 (Former) IT Consultant Confesses to SCADA Tampering

Hackers and malware writers and distributors have also attached power systems.[15] Cases summarized in the reference include

- 2000 Hacker Shocks Electric Company
- 2003 Slammer Worm Crashes Ohio Nuclear Plant Network
- 2006 National Nuclear Security Administration Computers Hacked; Info on 1500 Taken
- 2010 Stuxnet Worm Attacks SCADA Vulnerabilities

The Stuxnet incident is a significant development in information warfare. The European Network and Information Security Agency (ENISA) published a detailed analysis of the case in 2010:

> Stuxnet is a specialised malware targeting SCADA systems running Siemens SIMATIC® WinCC or SIMATIC® Siemens STEP 7 software for process visualisation and system control. SCADA in general refers to computer systems that monitor and control industrial processes, such as, e.g., those in nuclear power plants, or in facilities for water treatment.
>
> This highly sophisticated malware uses several vulnerabilities in the underlying Windows® operating system for infection and propagation. Infection works via USB-drives or open network shares. A root kit component hides the content of the malware on infected WinCC systems. An infected system can usually be controlled remotely by the attacker. In the end this means that the attacker has full control of the respective facility.[16]

ENISA also published a report on securing computer-controlled energy-distribution systems (*smart grids*) in December 2012.[17]

14.3.6 Transportation. Airplanes, trains, trucks, and ships are all likely targets for physical and information warfare. Because all of them are necessary to support the infrastructure by transporting personnel and materials, any disruption can cause severe problems. Because all of these transportation systems increasingly rely on sophisticated telecommunications and computing resources, they are subject to information warfare.

14.3.6.1 Aviation. The most visible, and potentially the most vulnerable, component of the transportation infrastructure is the aviation industry. Unlike the fly-by-the-seat-of-your-pants technology of aviation's early days, today's airplanes and the systems that dispatch and control them in flight are almost totally dependent on electronic communications and instruments, both analog and digital.

To a great extent, almost every airplane depends on its global positioning system (GPS) to determine its position in space, its course, speed, bearing to an airfield, and other important functions. Airplanes generally are required to fly at certain altitudes, in specific corridors, avoiding restricted areas, bad weather, and other aircraft. These requirements are met by a combination of GPS, ground and airborne radar, internal instruments, and communications from ground controllers. In the original design of these types of equipment, little or no consideration was given to security; as a result, all of them are susceptible to information warfare attacks.

The accuracy and reliability of GPS and airborne radar, however, has led federal aviation authorities to consider implementing a system wherein ground controllers and published restrictions would no longer determine altitude, speed, clearance distances, and other flight parameters. Instead, pilots would have the option to choose any flight parameter that they believed to be safe. This new system is intended to increase the number of flights that can safely traverse the limited airspace. It is undoubtedly capable of doing so, but at the same time, it will greatly increase the dangers of flight should information warfare be waged against airplanes and the aviation infrastructure.

14.3.6.2 Railroads. Less so than airplanes, but not to a negligible degree, trains are possible targets of IW. Train movements; switch settings, communications between engineers, trainmen, and control centers are all carried on by insecure radio communications and wired lines. Attacks against any or all of these can prevent the railroads from carrying out their important functions, possibly by causing disastrous wrecks.

14.3.6.3 Trucking. The great majority of domestic goods shipments are carried by tractor-trailer trucks. Foodstuffs, especially, depend on this relatively fast, reliable means of transportation. If even a short disruption were to be caused by IW, untold quantities of foodstuffs would rot in the fields, as would additional stockpiles awaiting distribution from central warehouses. Data for scheduling, routing, locating trucks, setting times and locations of pickup and delivery, and performing maintenance could be prevented from reaching their destinations.

14.3.6.4 Shipping. Ships are indispensable means for transporting vast quantities of materials over long distances. Navigational data, such as position, speed, course to steer, and estimated time of arrival, are a few of the parameters determined by computers and GPS on virtually every ship afloat. Conventional radar and communications by VHF and high-frequency radio are in common use, with satellite communications becoming more prevalent, despite an early start that met with technical and economic difficulties.

Radar and communications jamming are old established weapons of IW, as is interception of critical information. Little attention has been paid to security in designing or operating this equipment, and that places ships at great risk, as does the threat of physical attacks.

14.3.6.5 Other Transportation Vulnerabilities. Recognizing the importance of transportation to a nation's infrastructure, IW attackers could create wide-ranging disruptions if they were to intercept and successfully prevent receipt of critical information within the transportation industry. Recently, as a leader in new technology, the Port Authority of New York and New Jersey has begun converting to a wireless infrastructure at its many airports, train stations, bus terminals, tunnels, bridges, and shipping facilities. It requires no stretch of the imagination to predict what a determined attacker might accomplish in damaging or destroying such an infrastructure. The danger is especially great in light of the general lack of security from which wireless transmissions suffer.

When the World Trade Center (WTC) was destroyed by terrorist action, the Port Authority's offices in the WTC were completely destroyed, and more than 70 of its employees were officially listed as deceased or missing. Although that catastrophe pointed up the need for greater physical security, it also demonstrated how the Internet can be used in emergency situations. The Port Authority site, www.panynj.gov, was used to convey operational messages to the public as well as information for tenants, employees and prospective employees, vendors, suppliers, contractors, and the media.

14.3.7 Commerce. In 1924, in an address to the American Society of Newspaper Editors, President Calvin Coolidge said: "After all, the chief business of the American people is business. They are profoundly concerned with producing, buying, selling, investing, and prospering in the world. I am strongly of the opinion that the great majority of people will always find these are moving impulses of our life. …"[18]

Now, 90 years later at the time of writing, these statements are no less true. Producing, buying, selling, and investing are the commercial means by which U.S. citizens and guest workers can hope to achieve prosperity. Although not recognized earlier, infrastructure is the glue that ties these functions together and permits them to operate efficiently and economically.

If these bonds were to be broken, American business would come to a virtual standstill; it is that reality which makes the commercial infrastructure so inviting a target. Without complete, accurate, and current information, no investors would put their money at risk, and no transactions would take place among producers, buyers, and sellers.

In a populace lacking food, utilities, prescription drugs, money, and other necessities, civil disorder would be widespread. With the breakdown of commerce and the citizenry's unwillingness or inability to perform their customary functions, government at every level might cease to operate. This, in turn, would make military defensive actions highly problematic, and an enemy that combined IW with conventional force attacks would be difficult to resist.

On a less catastrophic level, there have been several cases of deliberate stock manipulation by means of insertion of false information into various news channels; an enemy could cause significant disruption in the stock market by forcing a few key stocks into unwarranted declines. In addition, the widespread use of automated trading tools that respond to significant drops in specific shares or in particular aggregate stock indexes could precipitate major economic problems in the developed world.

14.3.8 Financial Disruptions. Money is the lifeblood of every developed nation. For an economy to be healthy, its money supply, like the body's blood supply, must be strong, healthy, and free flowing. For an IW attacker, disruptions in the enemy's money supply and in its free flow are important objectives. Likely targets in the financial infrastructure include payment systems, investment mechanisms, and banking facilities.

14.3.8.1 Payment Systems. Every government employee, every member of the armed forces, every office worker, factory hand, service worker, engineer, and retired person—in fact, almost every individual in the United States—depends on regular receipt of funds necessary for survival. Paychecks, dividends, welfare and unemployment benefits, commissions, payments for products, and fees for services comprise most of the hundreds of millions of daily checks, direct deposits, and wire transfers without which most people would be unable to purchase their essential needs—assuming that products and services were available to meet those needs.

The great majority of payroll systems are computerized. Many of them, including those of the federal and state governments, depend on a few centralized computer payroll services. Even if those services were not damaged by infrastructure attacks, the banks on which payroll funds are drawn might be. This would halt, or at least impede, the cutting of physical checks, the direct deposits, cash withdrawals, wire transfers, and any other means by which payments are made. Such a situation has never occurred within the United States except in small local areas and for only brief periods of time. No one can predict what the consequences would be for a widespread attack, and surely no one would want to find out.

For more on banking payment systems, see Section 14.3.8.3.

14.3.8.2 Investment Mechanisms. Various stock, bond, and commodity exchanges provide the principal means by which individual, institutional, and corporate entities easily and expeditiously can invest in financial instruments and commodity goods.

With few exceptions, each exchange has all of its computers and communications located within a single facility, with connections to tens of thousands of terminals

worldwide. Disruption in these systems would not have as disastrous an effect as would a payment system disruption, but it would not be long before a breakdown in investment mechanisms would produce a commercial meltdown.

Because of the vast sums of money involved, exchange systems largely have been hardened against intrusion, and some have remote, redundant facilities, but there have been instances where hardware and software problems as well as physical exploits have brought down an exchange infrastructure.

14.3.8.3 Banking. The banking industry is the foundation of the modern financial system and, by extension, both American and foreign capitalist economies. At some point, every important financial transaction is conducted through the banking system. As such it is vital to economic health. With the advent of information warfare, the electronic, interdependent nature of banking—and finance in general—combined with its critical nature, makes the banking system a likely target for a strategic attack against a country. This is a new viewpoint for an industry focused on crime, traditional financial crises, and the more recent phenomenon of low-level hacking. It is critical, however, that we master this viewpoint and adapt our banking industry to it, for the threats information warfare poses are different from traditional bank security threats and will increase as the age of information warfare develops. Focused correctly, a well-prepared attack could cause chaos throughout the international system.[19]

The ubiquitous banking system is as highly automated and as security conscious as any element of the world's infrastructure. With ATMs, online banking, funds-transfer networks, and check clearing, banks are integral to virtually every commercial transaction.

As an example of the scope of banking operations involving money transfers, FEDWIRE, operated by the Federal Reserve Board, serves approximately 9,000 depository institutions as of January 2012, providing transfers that are immediate, final, and irrevocable.[20] In 2011, it processed 127 million transactions with a total value in excess of $663 trillion. The average daily volume in 2011 was over 506,000 transactions valued at more than $2.6 trillion.[21]

The Clearing House Interbank Payment System (CHIPS) "is responsible for over 95% of USD cross-border transactions, and nearly half of all domestic wire transactions, totaling $1.5 trillion daily [in 2012]."[22]

The Society for Worldwide Interbank Funds Transfer (SWIFT) "is a member-owned cooperative through which the financial world conducts its business operations with speed, certainty, and confidence. More than 10,000 financial institutions and corporations in 212 countries trust us every day to exchange millions of standardized financial messages. This activity involves the secure exchange of proprietary data while ensuring its confidentiality and integrity."[23] Information about dollar value is not made public, but the amounts are known to be huge and the traffic enormous. For example, by the end of September 2012, the annualized number of messages transferred for the year was over 339 million and the annualized volume of files transferred among institutions was over 22 million.[24]

If any of these systems were to be attacked successfully, the consequences for the financial well-being of many nations would be disastrous. Despite intensive efforts to safeguard the networks, attacks could be launched against the central computers, the computers of each user, and the networks that connect them.

14.3.9 Medical Security. In hospitals, as in group and private medical practice, the primary functions are carried out in a decentralized mode, making large-scale

attacks impracticable. However, ancillary functions, such as sending invoices to the government, to health maintenance organizations, and to individuals, for services provided, and placing orders for drugs and supplies, all require interconnections with centralized computers.

Although the medical profession is often slow to adopt new infrastructure elements, network-connected computers have been mandated at least for payments, and they are becoming increasingly popular for maintaining patient data, for research, and for other functions. There have been reports that hospital systems have been penetrated, with prescriptions switched and HIV-negative patients advised that their test results were positive.

See Chapter 71 of this *Handbook* for more detail about medical information security.

14.3.10 Law Enforcement. The objectives of law enforcement are to facilitate the apprehension of criminals and wrongdoers. To accomplish this, facilities in common use include computers in every squad car connected to precinct headquarters and networks that interconnect local, state, federal, and international databases. With local law enforcement, it is clear that jamming, or noise interference on emergency channels, or denial of computer services would greatly exacerbate the effects of physical attacks. At worst, a state of panic and chaos might ensue.

Another attack on law enforcement could be flash crowds—groups of people gathered into a single physical location through instructions sent electronically.[25] One commentator wrote in 2004,

> … [T]raining people to assemble on command in large numbers at, say, shoe stores, piano showrooms or restaurants for no good reason other than the fun of being part of a huge crowd is a perfect setup for creating an army of willing, mindless drones who will congregate on command at the site of a terrorist attack or at places where their presence will interfere with response to criminal or terrorist activities. Want to rob a bank in peace and quiet? Set up a conflict between two instant crowds to draw the police to an instant riot.[26]

14.3.11 International and Corporate Espionage. Espionage has been a recognized military activity since at least the biblical story of Joshua, one of 12 spies sent to explore the land of Canaan.[27] However, its application to civilian commerce dates only from the Industrial Revolution. Since then, industries and indeed nations have prospered to the extent that they could devise and retain trade secrets. In the United States, the unauthorized appropriation of military secrets has been legally proscribed since the country's inception, with penalties as severe as death, during wartime.

Only recently have economic espionage and the theft of trade secrets become the subjects of law, with severe penalties whether the law is broken within or outside of the United States or even via the Internet.

The Economic Espionage Act of 1996 was signed into law by President Clinton on October 11, 1996. Section 1832 provides that:

> **(A)** Whoever, with intent to convert a trade secret, that is related to or included in a product that is produced for or placed in interstate or foreign commerce, to the economic benefit of anyone other than the owner thereof, and intending or knowing that the offense will injure any owner of that trade secret, knowingly—
>
> **(1)** Steals, or without authorization appropriates, takes, carries away, or conceals, or by fraud, artifice, or deception obtains such information;
>
> **(2)** Without authorization copies, duplicates, sketches, draws, photographs, downloads, uploads, alters, destroys, photocopies, replicates, transmits, delivers, sends, mails, communicates, or conveys such information;

> **(3)** Receives, buys, or possesses such information, knowing the same to have been stolen or appropriated, obtained, or converted without authorization;
>
> **(4)** Attempts to commit any offense described in any of paragraphs (1) through (3); or
>
> **(5)** Conspires with one or more other persons to commit any offense described in any of paragraphs (1) through (3), and one or more of such persons do any act to effect the object of the conspiracy,
>
> Shall, except as provided in subsection (b), be fined under this title or imprisoned not more than 10 years, or both. (b) Any organization that commits any offense described in subsection (a) shall be fined not more than $5,000,000.[28]

Although the foregoing lists all of the actions that are proscribed, it is not specific as to which assets are to be protected as trade secrets. For this, see the Defense Security Service paper, "What Are We Protecting?"[29] There, the five basic categories of People, Activities/Operations, Information, Facilities, and Equipment/Materials are expanded into 42 specific assets, with the admonition that every company official must clearly identify to employees what classified or proprietary information requires protection. Only if the owner has taken reasonable measures to keep such information secret, and the information derives actual or potential economic value from not being generally known to or readily obtainable through proper means, will the courts view it as a trade secret.

For further information on intellectual property, including trade secrets, see Chapters 11 and 42 in this *Handbook*.

14.3.12 Communications. Communications are the means by which all elements of a civilization are tied together. Any significant destruction of communications media would disrupt the most important segments of society. Without adequate communications, transactions and services would come to a complete halt. In the United States, communications have been disrupted frequently, but fortunately, the infrastructure has been so vast and so diverse that the consequences have rarely been more than temporary. Even after the World Trade Center disaster of September 11, 2001, when Verizon's downtown telephone facilities centers were heavily damaged, service was restored within four days to the New York Stock Exchange and to other important users in the area.

Contrary to popular belief, the Internet is so widely used and concentrated in so few backbone points that a coordinated attack actually could destroy its functioning. For many years, backup facilities have included redundant computers and all of their associated peripherals, often in remote locations. Too often, however, alternate communications facilities are not provided. Unless this is rectified, the same disaster that brings down one installation could disable all.

14.3.13 Destabilization of Economic Infrastructure. A major difference between wealthy, developed nations and poor, undeveloped countries lies in the strength of their economic infrastructures. The existence of strong capital markets, stable banking and lending facilities, and efficient payment processes, all tied together by fast, technically advanced communications capabilities, is essential to healthy, growing economies.

At opposite ends of this spectrum lie Afghanistan and the United States. The perpetrators of the attacks on the World Trade Center and the Pentagon, identified as Osama bin Laden and his Al-Qaeda organization, operating out of Afghanistan, chose as their targets the symbols and the operating centers of America's military operations and of its economic infrastructure.

The recession of the late 2000s and early 2010s has illustrated the vulnerability of global economic systems to disruption. With hundreds of thousands thrown out of work and with investment capital drying up, the entire economic infrastructure of the world has already suffered a great blow without direct cyberattacks. Every effort must be bent toward preventing attacks that imperil the economic infrastructure of the world. Security can no longer be the duty of a few technical people; it has become everyone's responsibility.

14.4 SOURCES OF THREATS AND ATTACKS. The actual and potential originators of information warfare are numerous and powerful. One need not be paranoid to feel that an attack may come from any direction. This section lists sources that have already proven their capabilities for conducting cyberwar.

14.4.1 Nation-States. U.S. military preparations for cyberwar have been described in Section 14.3.2. This section details some of the measures that another great power—China—is effecting toward the same ends. Most of the material is from a paper entitled "Like Adding Wings to the Tiger: Chinese Information War Theory and Practice."[30]

14.4.1.1 China and Information Warfare. Although China is a nuclear power, it does not yet have the arsenal necessary to threaten a superpower like the United States. However, it can do so with its IW forces; adding wings to the tiger makes it more combat worthy. Nor is Chinese IW entirely theoretical. On August 3, 2000, the *Washington Times* reported that hackers suspected of working for a Chinese government institute took large amounts of unclassified but sensitive information from a Los Alamos computer system. A spokesman stated that "an enormous amount of Chinese activity hitting our green, open sites" occurs continuously.[31]

According to an article in the Chinese Armed Forces newspaper, the *Liberation Army Daily,* their first attack objectives will be the computer networking systems that link a country's political, economic, and military installations, as well as their general society.[32] A further objective will be to control the enemy's decision-making capability in order to hinder coordinated actions.

Expanding on Mao Zedung's theory of a People's War, IW can be "carried out by hundreds of millions of people using open-type modern information system."[33] In this war, combatants can be soldiers or teenagers, or anyone who has a computer as a weapon.[34] Ironically, China, with its long-standing fear of outside information as a possible spur to counterrevolutionary action, now views arming large numbers of intelligent people with computers and access to the Internet as a necessary survival measure. It remains to be seen just how many personal computers will be made available and how China will ensure that they will be used only as the government intends.

The "Annual Report to Congress on the Military Power of the People's Republic of China" from the U.S. Department of Defense has been issued every year since 2002. Reading through all the reports provides valuable perspective on the DoD view of information warfare capabilities of the People's Republic of China (PRC) and the People's Liberation Army (PLA). The 2011 edition included the following analysis:

- Cyberwarfare Capabilities. In 2010, numerous computer systems around the world, including those owned by the U.S. Government, were the target of intrusions, some of which appear to have originated within the PRC. These intrusions were focused on exfiltrating information. Although this alone is a serious concern, the accesses and skills required for these intrusions are similar to those necessary to conduct computer network attacks. China's 2010 *Defense*

White Paper notes China's own concern over foreign cyberwarfare efforts and highlighted the importance of cybersecurity in China's national defense.

- Cyberwarfare capabilities could serve PRC military operations in three key areas. First and foremost, they allow data collection through exfiltration. Second, they can be employed to constrain an adversary's actions or slow response time by targeting network-based logistics, communications, and commercial activities. Third, they can serve as a force multiplier when coupled with kinetic attacks during times of crisis or conflict.
- Developing capabilities for cyberwarfare is consistent with authoritative PLA military writings. Two military doctrinal writings, *Science of Strategy* and *Science of Campaigns*, identify information warfare (IW) as integral to achieving information superiority and an effective means for countering a stronger foe. Although neither document identifies the specific criteria for employing computer network attack against an adversary, both advocate developing capabilities to compete in this medium.
- The *Science of Strategy* and *Science of Campaigns* detail the effectiveness of IW and computer network operations in conflicts and advocate targeting adversary command and control and logistics networks to impact their ability to operate during the early stages of conflict. As the *Science of Strategy* explains,

  –In the information war, the command and control system is the heart of information collection, control, and application on the battlefield. It is also the nerve center of the entire battlefield. [Emphasis ours.][35]

In parallel with its military preparations, China has increased diplomatic engagement and advocacy in multilateral and international forums where cyberissues are discussed and debated. Beijing's agenda is frequently in line with the Russian Federation's efforts to promote more international control over cyberactivities. China has not yet agreed with the U.S. position that existing mechanisms, such as International Humanitarian Law and the Law of Armed Conflict, apply in cyberspace. China's thinking in this area is evolving as it becomes more engaged."

14.4.1.2 Strategies. The People's Liberation Army (PLA) with 1.5 million reserve troops has been carrying out IW exercises on a wide scale. One such exercise, in Xian Province, concentrated on conducting information reconnaissance, changing network data, releasing information bombs, dumping information garbage, disseminating propaganda, applying information deception, releasing clone information, organizing information defense, and establishing spy stations.[36] The antecedents of these tactics can be found in a book of unknown authorship, first mentioned about 1,500 years ago, entitled *The Secret Art of War: The 36 Stratagems.* Strategy 25 advises:

> *Replace the Beams with Rotten Timbers.* Disrupt the enemy's formations, interfere with their methods of operations, change the rules which they are used to following, and go contrary to their standard training. In this way you remove the supporting pillar, the common link that makes a group of men an effective fighting force.[37]

The 36 stratagems deserve close study; many of them are obviously in use even today by China and others. For example, strategy 3 says:

> *Kill with a Borrowed Sword.* When you do not have the means to attack your enemy directly, then attack using the strength of another.

Lacking the weapons to attack the United States directly, the perpetrators of the WTC attack used the airliners belonging to their targets.

Strategy 5 says:

> *Loot a Burning House.* When a country is beset by internal conflicts, when disease and famine ravage the population, when corruption and crime are rampant, then it will be unable to deal with an outside threat. This is the time to attack.

Some of the strategies might well be employed by the United States. For example, strategy 33 advises:

> *The Strategy of Sowing Discord*. Undermine your enemy's ability to fight by secretly causing discord between him and his friends, allies, advisors, family, commanders, soldiers, and population. While he is preoccupied settling internal disputes his ability to attack or defend, is compromised.

To accomplish this, IW may prove to be an effective weapon.

14.4.1.3 Training. Several high-level academies and universities have been established to conduct IW instruction for the PLA. In addition, training is planned for large numbers of individuals to include:

- Basic theory, including computer basics and application, communications network technology, the information highway, and digitized units
- Electronic countermeasures, radar technology
- IW rules and regulations
- IW strategy and tactics
- Theater and strategic IW
- Information systems, including gathering, handling, disseminating, and using information
- Combat command, monitoring, decision making, and control systems
- Information weapons, including concepts, principles of soft and hard destruction, and how to apply these weapons
- Simulated IW, protection of information systems, computer virus attacks and counterattacks, and jamming and counterjamming of communications networks[38]

It is doubtful that all of these training objectives have been accomplished, but there seems to be a major commitment to do so, and sooner rather than later.

China and the United States are only two of the nations that are openly preparing for, and actually engaged in, information warfare. It is obvious that many others are similarly involved and that these measures, combined with conventional weapons or weapons of mass destruction, have the potential to elevate warfare to a destructive level never before possible and hardly conceivable.

14.4.2 Cyberterrorists

> "Cyberterrorism" means intentional use or threat of use, without legally recognized authority, of violence, disruption, or interference against cybersystems, when it is likely that such use would result in death or injury of a person or persons, substantial damage to physical property, civil disorder, or significant economic harm.[39]

Cyberterrorists, those who engage in cyberterrorism, generally are able to carry out the same sort of cyberwar as nation-states; in fact, they may be state-sponsored. The major difference is that terrorist attacks are usually hit-and-run, where nations are capable of sustained and continuous operations. Although conventional warfare always was carried out in an overt fashion, it is the nature of IW that it can be engaged in without a declaration of war and without any clear indication of who the attacker actually is. In fact, it may not be recognized that a war is being conducted; it may

seem only that a series of unfortunate, unconnected natural failures of computers and communications are disrupting an economy.

Terrorists, especially when state-sponsored, would be very likely to conceal their IW activities in this manner, so as to avoid the retribution that would inevitably follow. However, some terrorists would publicly take credit for their actions, in order to bolster their apparent strength and to gather added support from like-minded individuals and organizations.

The seriousness of terrorist threats after 9/11 resulted in Executive Order 13228 of October 8, 2001, establishing the Office of Homeland Security and the Homeland Security Council.[40] The mission of the Office was to "develop and coordinate the implementation of a comprehensive national strategy to secure the United States from terrorist threats or attacks." Its function was "to coordinate the executive branch's efforts to detect, prepare for, prevent, protect against, respond to, and recover from terrorist attacks within the United States."

The Department of Homeland Security was mandated by Congress on January 24, 2003, and was fully formed on March 1, 2003. Celebrating its tenth anniversary in 2013, the department employs more than 200,000 people dedicated to fulfilling its mission.

On February 15, 2005, Michael Chertoff was sworn in as the second secretary. His five goals:

1. Protect our Nation from Dangerous People
2. Protect our Nation from Dangerous Goods
3. Protect Critical Infrastructure
4. Strengthen our Nation's Preparedness and Emergency Response Capabilities
5. Strengthen and Unify Operations and Management[41 26]

On April 30, 2008, Secretary Chertoff, recognizing new realities, said:

> [T]he technology of the 21st Century is changing so rapidly that many of our rules and procedures, which were built at a time that we had a certain kind of communication system and a certain kind of analog set of processes, that legal structure seems woefully inadequate to a digital age when the movement of communications is not rooted in any one place and when it's very difficult to take the concepts which made a lot of sense in the days of the rotary telephone and apply them in the world of voice over internet protocols.[42]

In March 2010, a report on comments by FBI Director Robert Mueller and former White House "terrorism czar" Richard Clarke included this summary:

> "As you well know, a cyber-attack could have the same impact as a well-placed bomb," Mueller said. "In the past 10 years, Al-Qaeda's online presence has become as potent as its in-world presence."
>
> Al-Qaeda uses for the Internet range from recruiting members and inciting violence to posting ways to make bio-weapons and forming social-networks for aspiring terrorists, according to Mueller. "The cyber-terrorism threat is real and rapidly expanding," Mueller said. "Terrorists have shown a clear interest in hacking skills and combining real attacks with cyber attacks."
>
> Threats are also rising from online espionage, with hackers out for source code, money, trade, and government secrets, according to the FBI. "Every major company in the U.S. and Europe has been penetrated—it's industrial warfare," said Richard Clarke, who was a White House adviser under three prior U.S. presidents. "All the little cyber-devices that the companies here sell have been unable to stop them. China and Russia are stealing petabytes of information."

> Clarke, now a partner at Good Harbor Consulting firm, was among the RSA panelists discussing cyber-warfare. "Nation states have created cyber-warfare units. They are preparing the battlefield," Clarke said. "We have the governments of China and Russia engaging in daily activities successfully that the U.S. government and private industry are not stopping and they are stealing anything worth stealing. …
>
> Mueller urged computer security professionals to join in a united, international alliance with law enforcement agencies to battle enemies in cyberspace. He credited such teamwork with resulting in the recent arrest of three men in Spain suspected of running a network of nearly 13 million computers secretly infected with malicious software and used for nefarious deeds.
>
> Mueller called on victims of cyber-attacks to break the pattern of remaining silent out of fear that reporting crimes would hurt their positions in the marketplace. "Maintaining the code of silence will not benefit you or your clients in the long run," Mueller said. "We must continue to do everything we can together to minimize and stop these attacks."[43]

The challenges faced by the Department of Homeland Security are multitudinous and complex. Whether it proves effective in reducing or eliminating terrorism within the United States will depend on solving the problems of overlapping authorities, inertia, incompatible databases, turf wars, funding, management, the predictability of terrorist actions, and a host of political and technological issues.

14.4.3 Corporations. The threats aimed at or directed by corporations are far less deadly than those of the military or of terrorists, but they are no less pervasive. Thefts of data, denial of service, viruses, and natural disasters traditionally have been at the heart of individual corporate security concerns. These concerns have not abated, but to them have been added fears that attacks on large segments of the information infrastructure are more likely to create damage than is an attack against any single enterprise. To guard against this, every installation should operate behind strong firewalls and effective access controls.

In the wake of the September 11 attacks, Richard Clarke, who had been National Coordinator for Security, Infrastructure Protection, and Counterterrorism since May 1998, was appointed to a new post. As special advisor to the president for cyberspace security, Mr. Clarke warned that terrorists are out to hurt our economy and that they can use viruses in massive, coordinated attacks against corporate IT systems. He recommended, at a minimum, that disaster recovery plans include near-online, offsite backup facilities and redundant communications paths.

14.5 WEAPONS OF CYBERWAR. The weapons used in information warfare have existed for many years, but newer and more malevolent versions are produced with increasing frequency. For this reason, system security cannot be considered as static, but rather as part of an ongoing process that must be continuously monitored and strengthened. This section briefly describes the most common and most dangerous IW weapons, with references to other chapters where more detailed information is available.

14.5.1 Denial of Service and Distributed Denial of Service. Denial of service (DoS) and distributed denial of service (DDoS) are means by which computers, network servers, and telecommunications circuits can be partially or completely prevented from performing their designated functions. Any computer element that has been designed for a specific maximum capacity, if flooded by messages or data inputs that greatly exceed that number, can be slowed or even brought to a complete halt.

A DoS attack is carried out by a single computer that has been programmed to overwhelm the target system's capacity, usually by generating, automatically, a very large number of messages. A DDoS attack is implemented by planting a small program on hundreds or thousands of unaware computers. At a signal from the attacker, all of the agents (sometimes called zombies or daemons) are caused to send many messages simultaneously, thus flooding the victim's system or preempting all of its bandwidth capacity.

On April 26, 2007, a page-one article in the *New York Times* reported on what some Estonian authorities described as the first war in cyberspace. It was precipitated by the removal from a park in Tallinn of a bronze memorial to the Soviet soldiers of World War II. It was believed, but not proven, that the Russian government, or individual activists, had used DDoS attacks to bring down computers propagating the Websites of the Estonian president, prime minister, and Parliament as well as of banks and newspapers. The attacks were finally brought under control with the help of experts from NATO, the European Union, the United States, Finland, Germany, Slovenia, and Israel. Details of many DoS and DDos attacks and the recommended defenses are contained in Chapter 18 of this *Handbook.*

14.5.2 Malicious Code. Malicious code includes viruses, worms, and Trojan horses, as described in Chapter 16. Mobile code, such as Java, ActiveX, and VBScript, was developed to increase the functionality of Websites, but all three, as described in Chapter 17, also can be used maliciously.

There have been innumerable instances where malicious code has been used to damage or deface Websites, both civilian and military. Apparently, all of these exploits have been perpetrated by single individuals or by very small groups of unaffiliated crackers. However, in the event of actual cyberwar, it seems certain that large groups of coordinated, technically knowledgeable attackers will attempt to wreak havoc on their opponents' infrastructures through the use of malicious code.

In 2012, news reports indicated that Flame, malware with a relationship to Stuxnet, may, similarly, have been developed by U.S. and Israeli cyberwarfare specialists. Flame was not used for sabotage but rather for data theft.[44] It was held to be responsible for stealing Iranian passwords, network descriptors, and even data files. Armed with this information, Stuxnet was allegedly used to destroy Iranian centrifuges enriching uranium as a component of atomic weapons. It accomplished this by controlling the controllers that set the speed of the centrifuges. When the speeds were set excessively high, they damaged or destroyed the centrifuges.

Because most of the hardware elements of these systems are commercially available, and because the software is readily duplicated, the threats to military and commercial applications are imminently capable of extensive and costly attacks. Just as U.S. military and governmental agencies, and most of their allies, are engaged in large-scale operations to develop defensive capabilities, it is essential that all commercial enterprises exert major efforts to do the same. Initiatives have begun to form close working relationships between government and the private sector. Also, industry groups have begun advocating relaxation of those laws that prohibit close cooperation between competitors. This will be necessary before information can be shared as required to strengthen the infrastructure. Similarly, groups are requesting that shared information be protected from those who would use the Freedom of Information Act to force disclosure.

Every prudent organization will support these initiatives and will work with appropriate government agencies and industry groups to ensure its own survival and the welfare of the country itself.

14.5.3 Cryptography. Military operations, since the earliest recorded times, have utilized cryptography to prevent critical information from falling into enemy hands. Today, information is a vastly more important resource than ever before, and the need for cryptography has increased almost beyond measure. Not only the military, but indeed every financial institution, every competitive commercial enterprise, and even many individuals feel impelled to safeguard their own vital information. At the same time, access to the secret information of enemies and opponents would provide inestimable advantages.

Recognizing this, powerful supercomputers, directed by mathematicians, theoretical scientists, and cryptographers, are being applied to improving the processes of encryption and decryption. The most notable achievement in the recent past was the British construction of a computerized device to break the German Enigma code. The information thus obtained has been widely credited with a significant role in the outcome of World War II.

The development of effective mechanisms for spreading computations over millions of personal computers has greatly reduced the time required for brute force cracking of specific encrypted messages; for example, messages encrypted using the 56-bit Digital Encryption Standard (DES) were decrypted in four months using 10,000 computers in 1997, 56 hours using 1,500 special-purpose processors in 1998, and 22 hours using 100,000 processors in 1999.[45]

A major issue, yet to be resolved, is the strength of cryptographic tools that may be sold domestically or exported overseas. The contending forces include producers of cryptographic tools who believe that if the strength of their product is in any way restricted, they will lose their markets to producers in other countries with more liberal policies. Similarly, proponents of privacy rights believe that unbreakable cryptographic tools should be freely available.

The countervailing view is that virtually unbreakable cryptographic tools shipped overseas will inevitably find their way into the hands of unfriendly governments, which may use them in conducting cyberwars against us. Domestically, law enforcement agencies believe that they should have "back-door" entry into all cryptographic algorithms, so that they may prevent crimes as wide-ranging as embezzlement, drug trafficking, and terrorism.

As domestic crimes and terrorist attacks grow in number and intensity, it seems certain that at least a few civil liberties, including privacy rights, may be infringed. The hope is that an optimum balance will be struck between the need for security and the core values of our democracy.

For more on privacy in cyberspace, see Chapter 69 in this *Handbook.*

14.5.4 Psychological Operations. Psychological operations (PSYOP) may be defined as planned psychological activities directed to enemy, friendly, and neutral audiences in order to influence their emotions, motives, attitudes, objective reasoning, and behaviors in ways favorable to the originator. The target audiences include governments, organizations, groups, and individuals, both military and civilian.

One of the most potent weapons in information warfare, PSYOP attempts to:

- Reduce morale and combat efficiency within the enemy's ranks
- Promote mass dissension within, and defections from, enemy combat units and/or revolutionary cadres

- Support our own and allied forces cover and deception operations
- Promote cooperation, unity, and morale within one's own and allied units, as well as within friendly resistance forces behind enemy lines[46]

The information that accomplishes these ends is conveyed via any media: by printed material such as pamphlets, posters, newspapers, books, and magazines, and by radio, television, personal contact, public address systems, and of increasing importance, through the Internet.

A classic example of successful PSYOP application was the deception practiced prior to the Allied invasion of the European mainland. Through clever "leaks," false information reached Germany that General Patton, America's most celebrated combat commander, was to lead an army group across the English Channel at Pas de Calais. As a consequence, German defensive forces were concentrated in that area. For weeks after the Normandy invasion was mounted, Hitler was convinced that it was just a feint, and he refused to permit the forces at Calais to be redeployed. Had this PSYOP failed, and had more of Germany's defensive forces been concentrated in Normandy, the Allied landing forces might well have been thrown back into the sea.

Although generally considered not to involve a PSYOP action, the September 11 attacks and the subsequent spread of anthrax spores made clear that a physical action can have the greatest and most far-reaching psychological effects. Beyond mourning the death of almost 3,000 innocent civilians, the new sense of vulnerability and powerlessness caused great psychological trauma throughout the nation and much of the Western world. The full consequences to the travel, entertainment, and hospitality industries, as well as to every segment of the world economy, are likely to be both disastrous and long-lasting.

A major, integrated, expert PSYOP mission to restore morale and encourage behavior can halt or reverse a downward spiral, but worldwide recessions and acts of nature, such as cyclones, hurricanes, and earthquakes, can do more than PSYOP actions to demoralize a nation.

14.5.5 Physical Attacks. Prior to September 11, 2001, physical attacks, as a part of cyberwar, were generally considered in the same light as attacks against any military objective, and defensive measures were instituted accordingly. In the civilian sector, starting with student attacks against academic computers in the 1960s and 1970s, there have been occasional reported physical attacks against information processing resources. Although access controls have been almost universally in place, their enforcement often has been less than strict.

Another indication of the susceptibility of the information infrastructure to physical attack is the prevalence of "backhoe attacks" in which construction crews accidentally slice through high-capacity optic cables used for telecommunications and as part of the Internet backbones.[47] The signs indicating where not to dig can serve as markers for those targeting single points of failure.

A related vulnerability is undersea telecommunications cables, which are unprotected against accidental—or deliberate—damage from ship anchors and from other objects or tools. Breaks in these cables can interrupt the Internet and telephone networks on a global scale.[48]

The destruction of the WTC and a portion of the Pentagon have brought the possibility of additional physical attacks very much into the forefront of cyberwar thinking, for both the military and the civilian infrastructures. Car bombings and packaged bombs

had become almost commonplace, especially in the Mideast. Successful attacks had been launched against U.S. embassies and troop barracks, as well as against Israel, England, Spain, and France. To guard against such actions, perimeter defenses were widened, and in some areas personal searches at strategic points were instituted.

These defenses have proven to be of limited value, and suicide bombers seem to be increasing in numbers and in the effectiveness of their weapons. The use of commercial aircraft, fully loaded with fuel, as manned, guided missiles was apparently never considered prior to 11 September. After that date, there has been widespread recognition that protective measures must be taken that will prevent a recurrence of those tragic events. Airport security has become a direct federal responsibility, under a new Transportation Security Administration in the Department of Transportation. On November 11, 2001, President Bush signed a bill that requires all airport baggage screeners to be U.S. citizens and to undergo criminal background checks before becoming federal employees. At many airports, security is provided by private contractors. The protective measures in common use are considered to be pointless, inconvenient, and ineffective by many travelers. Although even minimal safeguards against known weapons are being debated, there appears to be little thinking directed toward other types of attacks that might even now be in the planning stage.

14.5.6 Biological and Chemical Weapons and Weapons of Mass Destruction. Although the use of these weapons can affect every element of society, they have a particular potency in destroying the infrastructure of a targeted nation. The WTC attacks have had long-lasting psychological effects, but the results of the anthrax dissemination may be even more deeply traumatic. Already, the presence of anthrax spores has interfered with the functioning of the Congress, the Supreme Court, the U.S. Postal Service, hospitals, and other institutions. Although the furor over these attacks, as well as their incidence, has dissipated, there may be even more such attacks in the future. Unless any future culprit is apprehended quickly, and countermeasures taken immediately, damage to the infrastructure could be extensive.

14.5.7 Weapons Inadvertently Provided. There are many widespread vulnerabilities to computer systems that are not created as weapons, but whose presence makes the targets of cyberwar highly vulnerable. Poor software designs and inadequate quality control create opportunities for attackers to damage or destroy information, and the information systems themselves. Chapters 38 to 40 of this *Handbook* are especially useful in identifying and eliminating these sources of security vulnerabilities.

14.6 DEFENSES. A variety of defenses may be employed both to prevent attacks and to mitigate their effects. Because each of these defenses may have only limited utility, it is evident that new and more effective defenses must be developed.

14.6.1 Legal Defenses. As a defense against IW attacks or as a framework for apprehending and prosecuting attackers, the international legal system has been generally ineffective. The reasons for this include:

- Information warfare is not prohibited under the United Nations (UN) Charter, unless it directly results in death or property damage.
- Laws that are not recognized and enforced lose their power to compel actions.
- There is little or no police power to enforce those few laws that do exist.

- The issue of sovereignty as it relates to transborder communications is unresolved.
- Neither the United States nor any other major power has pressed for international laws to govern information warfare. This may be attributed to the fact that such laws, while desirable for defense, would impair the nation's own offensive operations.
- Many nations do not recognize cyberwar attacks as criminal actions.
- In many lands, political considerations determine judicial outcomes.
- Few countries support extradition of their citizens even when indicted for terrorist or criminal activities.
- Terrorists, drug cartels, the international mafia, and even individual hackers have every reason to circumvent the law, and usually possess the resources that enable them to do so.
- Identifying attackers may be difficult or even impossible.
- New technologies arrive at a rate much faster than appropriate legislation.

Further acting to constrain law as a deterrent is the fact that there has been no universal acceptance of definitions for IW-relevant terminology: Attacks, acts of war, aggression, hostilities, combatants, crimes, criminals—all remain vague concepts. Until such terms, as applied to IW, are clearly defined, there can be no legal strictures against them.

The difference between acceptable and unacceptable targets is obscured by the dual-use, civilian and military, characteristics of infosystems and infrastructures. Similarly, it is difficult to condemn denial of service, when peacetime boycotts and economic sanctions are widely applied to further economic or political ends.

Clearly, legal defenses against cyberwar are inadequate at this time. Whether the United States will pursue effective international legislation remains doubtful, until the question of building adequate defenses, without hobbling offensive operations, is resolved.

14.6.2 Forceful Defenses. If IW attacks are accepted as acts of war, the use of retaliatory military force would be highly likely. The strategic and tactical decisions that would follow are well beyond the scope of this chapter, but six considerations are relevant.

1. The United States is growing reluctant to engage in combat without the sanction of the United Nations and without the concurrence of major allies. If the provocation is limited to an IW attack, it may be difficult to build a coalition or even to avoid UN condemnation.
2. The identity of the attacker may be unclear. Even after the September 11 attacks, the United States had no enemy that admitted culpability. As a consequence, the United States could not declare war on any nation or state but could only declare a war on "terrorism."
3. The attacker may be misidentified. Through the use of "spoofing" and routing an attack through unaware nations, the anonymous culprit may escape detection, while blame falls on an innocent victim.
4. There may be difficulty in determining whether a particular event is an act of information warfare or simply the result of errors, accidents, or malfunctions.
5. The attackers may not be a foreign government, against whom war can be declared, but a criminal organization, a disaffected group, activists, commercial competitors, or even individuals bent on mischief.

6. The United Nations, and international sentiment in general, requires that military force only be used in response to armed attack and, further, that the response be proportional to the attack that provoked it.

In light of these considerations, it seems unlikely that information warfare, unless it results in catastrophic injuries and deaths, will be met by a forceful reaction. The top secret Presidential memo leaked on June 7, 2013, and discussed in Section 14.3.3 clearly indicates that armed actions may now be initiated without the impetus of injuries and deaths.

14.6.3 Technical Defenses. The technical defenses against IW are many and varied. Almost the entire contents of this volume are applicable to safeguarding against cyberwar attacks. These same measures can prove equally effective in defending against IW, criminals, activists, competitors, and hackers.

14.6.4 In-Kind Counterattacks. A cyberwar defense that has been used often is an in-kind counterattack, where flaming is met by flaming, DDoS by DDoS, site defacement by site defacement, and propaganda by propaganda. Recent examples include exchanges between Israelis and Arabs, Kashmiris and Indians, Serbs and Albanians, Indians and Pakistanis, Taiwanese and Chinese, and Chinese and Americans.

Although there may be personal satisfaction in originating or responding to such attacks, the net effect is usually a draw, and, therefore, in-kind attacks generally have been short-lived. In the future, such attacks may no longer be the output of only a few individuals, but may be mounted by large numbers of similarly minded cyberwarriors, organized into coordinated groups, with sophisticated tools and with covert or overt state sponsorship.

In that event, the asymmetric nature of the adversaries' infrastructures would be telling. Clearly, if the Taliban, for example, were to mount another full-scale cyberterrorist attack against the United States with the help of their supporters throughout the world, the effects could be devastating. Although the United States might mount a highly sophisticated in-kind response, it probably would have no effect on the Taliban's organization, its economy, its military effectiveness, or its ability to carry out suicide missions, biological warfare, or other physical attacks. A great and powerful nation may lack the ability to destroy a small, primitive, almost nonexistent infrastructure.

A serious problem with any kind of counterattack is that the origins of cyberattacks through the Internet using Internet Protocol version 4 (IPv4) are easily masked, or *spoofed*, because of the lack of mandatory, verifiable source authentication. It is possible that poorly analyzed data about the supposed attackers could lead to a counterattack against innocent victims.

14.6.5 Integration of Cyberwarfare into Military Planning. The United States Cyber Command (USCYBERCOM) was initiated in 2009 as a subset of the United States Strategic Command and became fully operational in its Fort Meade headquarters in 2010.[49] Its mission and focus are defined as follows:

- Mission: USCYBERCOM is responsible for planning, coordinating, integrating, synchronizing, and directing activities to operate and defend the Department of Defense information networks and when directed, conducts full-spectrum military cyberspace operations (in accordance with all applicable laws and regulations) in order to ensure U.S. and allied freedom of action in cyberspace, while denying the same to our adversaries.

- Focus: The command is charged with pulling together existing cyberspace resources, creating synergy that did not previously exist and synchronizing war-fighting effects to defend the information security environment.
- The Command centralizes direction of cyberspace operations, strengthens DoD cyberspace capabilities, and integrates and bolsters DoD's cyberexpertise. USCYBERCOM improves DoD's capabilities to ensure resilient, reliable information and communication networks, counter cyberspace threats, and assure access to cyberspace. The command works closely with interagency and international partners in executing the cybermission.

USCYBERCOM has initiated training exercises for U.S. military forces. The first *Cyber Flag* exercise was held in 2011 at Nellis Air Force Base; the second was in 2012.[50] The 2012 "… exercise saw approximately 700 participants, up from last year's 300, and doubled the network size. All participants had a specific role to play, playing the part of a U.S. team or role-playing an adversary."

In January 2013, plans surfaced for a major investment by the Pentagon in cybersecurity:

> The expansion would increase the Defense Department's Cyber Command by more than 4,000 people, up from the current 900, an American official said. Defense officials acknowledged that a formidable challenge in the growth of the command would be finding, training, and holding onto such a large number of qualified people.
>
> The Pentagon "is constantly looking to recruit, train and retain world class cyberpersonnel," a defense official said Sunday.
>
> "The threat is real and we need to react to it," said William J. Lynn III, a former deputy defense secretary who worked on the Pentagon's cybersecurity strategy.
>
> As part of the expansion, officials said the Pentagon was planning three different forces under Cyber Command: "national mission forces" to protect computer systems that support the nation's power grid and critical infrastructure; "combat mission forces" to plan and execute attacks on adversaries; and "cyberprotection forces" to secure the Pentagon's computer systems.[51]

14.6.6 Cooperative Efforts. Although the United States has been moderately successful in building coalitions in support of military operations, it has shown little inclination to build an international consensus dealing with information warfare. This may be so because of the legal difficulties outlined in Section 14.6.1 or because any prohibitions against offensive cyberwar will limit United States options. Nevertheless, whether by treaty, convention, agreement, or UN directive, technical people, diplomats, and statesmen of all well-intentioned countries should work together to define unacceptable and harmful actions and to devise means for detecting, identifying, and punishing those who transgress.

In June 2013, the North Atlantic Treaty Organization (NATO) representatives participated in a discussion of "how best to defend against cyber threats."[52] NATO posted this summary of its stance:

> Against the background of rapidly developing technology, NATO is advancing its efforts to confront the wide range of cyber threats targeting the Alliance's networks on a daily basis. NATO's Strategic Concept and the 2012 Chicago Summit Declaration recognise that the growing sophistication of cyber attacks makes the protection of the Alliance's information and communications systems an urgent task for NATO, and one on which its security now depends.
>
> In June 2011, NATO adopted a new cyber defence policy and the associated Action Plan, which sets out a clear vision of how the Alliance plans to bolster its cyber efforts. This policy

reiterates that any collective defence response is subject to decisions of the North Atlantic Council, NATO's principal political decision-making body.

The revised policy offers a coordinated approach to cyber defence across the Alliance. It focuses on the prevention of cyber attacks and building resilience. All NATO structures will be brought under centralised protection and NATO will enhance its capabilities to deal with the vast array of cyber threats it currently faces, including through integrating them into the NATO Defence Planning Process. This way Allies will ensure that appropriate cyber defence capabilities are included as part of their planning to protect information infrastructures that are critical for core Alliance tasks. The revised cyber defence policy also stipulates NATO's principles on cyber defence cooperation with partner countries, international organisations, the private sector and academia.[53]

14.7 SUMMARY. The potential for information warfare to damage or destroy the infrastructure of any nation, any corporation, or, in fact, any civilian, governmental, or military entity is unquestionable. Until now, the only incidents have been isolated and sporadic, but the possibility of sustained, coordinated, simultaneous attacks is strong. If these attacks are combined with physical, chemical, or biological warfare, the effects are certain to be devastating.

Although the types of potential attackers, and the probable weapons they will use, are well known, the available defenses do not at this time offer any great assurance that they will be effective. The United States and many of its allies are engaged in great efforts to remedy this situation, but formidable obstacles are yet to be overcome. The military is generally better prepared than the civilian sector, but much of the military's infrastructure is woven into and dependent on transportation, communications, utilities, food production and distribution, and other vital necessities that are owned by private enterprises.

Recent terrorist attacks and the probability of future offensives should serve as an immediate impetus to devote whatever resources are needed to combat the threats to our way of life and, in fact, to our very existence.

14.8 FURTHER READING

Armistead, E. L. *Information Operations: Warfare and the Hard Reality of Soft Power*. Potomac Books, 2004.

Armistead, E. L. *Information Warfare: Separating Hype from Reality*. Potomac Books, 2007.

Armistead, E. L. *Information Operations Matters: Best Practices*. Potomac Books, 2010.

Arquilla, J., and D. Ronfeldt, eds. *In Athena's Camp: Preparing for Conflict in the Information Age*. RAND Corporation, 1997. Available free in parts as PDF files from http://rand.org/pubs/monograph_reports/MR880/

Campen, A. D., and D. H. Dearth, eds. *Cyberwar 3.0: Human Factors in Information Operations and Future Conflict*. Fairfax, VA: AFCEA International Press, 2000.

Clarke, R. A. *Cyber War: The Next Threat to National Security and What to Do About It*. Ecco, 2010.

Cohen, F. *World War 3: We Are Losing It and Most of Us Didn't Even Know We Were Fighting in It—Information Warfare Basics*. Fred Cohen & Associates, 2006.

Denning, D. E. *Information Warfare and Security*. Addison-Wesley, 1998.

Erbschloe, M., and J. Vacca. *Information Warfare*. McGraw-Hill, 2001.

Greenberg, L., S. E. Goodman, and K. J. Soo Hoo. *Information Warfare and International Law*. National Defense University Press, 1998.

Hall, W. M. *Stray Voltage: War in the Information Age*. U.S. Naval Institute Press, 2003.

Henry, R., and C. E. Peartree, eds. *The Information Revolution and International Security*. Center for Strategic and International Studies, 1998.

Kahn, D. *The Codebreakers*. Scribner, 1996.

Kramer, F. D., S. H. Starr, and L. Wentz, eds. *Cyberpower and National Security*. Potomac Books, 2009.

Lesser, I. O., B. Hoffman, J. Arquilla, D. Ronfeldt, and M. Zanini. *Countering the New Terrorism*. RAND Project Air Force, 1999. Available free in parts as PDF files from http://rand.org/pubs/monograph_reports/MR989/

Macdonald, S. *Propaganda and Information Warfare in the Twenty-First Century: Altered Images and Deception Operations*. Routledge, 2007.

Marsh, R. T., chair. *Critical Foundations: Protecting America's Infrastructures. The Report of the President's Commission on Critical Infrastructure Protection*. 1997; www.fas.org/sgp/library/pccip.pdf

Price, A., and C. A. Horner. *War in the Fourth Dimension: U.S. Electronic Warfare, from the Vietnam War to the Present*. London, UK: Greenhill Books/Lionel Leventhal, 2001.

Rattray, G. J. *Strategic Warfare in Cyberspace*. MIT Press, 2001.

Reveron, D. S. *Cyberspace and National Security: Threats, Opportunities, and Power in a Virtual World*. Georgetown University Press, 2012.

Rid, T. *Cyber War Will Not Take Place.* Oxford University Press, 2013.

Rosenzweig, P. *Cyber Warfare: How Conflicts in Cyberspace Are Challenging America and Changing the World*. Praeger, 2013.

Schwartau, W. *Information Warfare: Chaos on the Electronic Superhighway*, 2nd ed. Thunder's Mouth Press/Perseus Publishing Group, 1996.

Zalmay, K., and J. P. White, eds. *Strategic Appraisal: The Changing Role of Information in Warfare*. McGraw-Hill, 1999.

14.9 NOTES

1. W. J. Clinton, "Critical Infrastructure Protection," Presidential Decision Directive 63, May 22, 1998, www.fas.org/irp/offdocs/pdd/pdd-63.htm
2. J. L. Brock, "Critical Infrastructure Protection: Fundamental Improvements Needed to Assure Security of Federal Operations," GAO/T-AIMD-00-7, Testimony before the Subcommittee on Technology, Terrorism and Government Information, Committee on the Judiciary, U.S. Senate, October 6, 1999, www.gao.gov/archive/2000/ai00007t.pdf
3. L. Wright, "Protecting the Homeland: Report of the Defense Science Board Task Force on Defensive Information Operations 2000 Summer Study, Vol. II." Office of the Undersecretary of Defense for Acquisition, Technology, and Logistics (March 2001), www.acq.osd.mil/dsb/reports/dio.pdf (URL inactive)
4. T. P. M. Barnett, "The Seven Deadly Sins of Network-Centric Warfare," United States Naval Institute Proceedings 125, No. 1 (January 1999): 36–39, www.usni.org/magazines/proceedings/1999-01/seven-deadly-sins-network-centric-warfare
5. G. G. Gilmore, "Navy-Marine Corps Intranet Girds for Cyber-Attacks," Armed Forces Press Service, July 6, 2001, www.defenselink.mil/news/newsarticle.aspx?id=44745

6. U.S. Office of the Secretary of Defense, "Annual Report to Congress: Military and Security Developments Involving the People's Republic of China 2013," www.defense.gov/pubs/2013_China_Report_FINAL.pdf
7. Joint Chiefs of Staff, "Joint Doctrine for Information Operations," Joint Publication 3-13, 2006, www.dtic.mil/doctrine/jel/new_pubs/jp3_13.pdf (URL inactive)
8. W. S. Cohen, Annual Report to the President and the Congress: Secretary of Defense, 2001, www.dod.mil/execsec/adr2001/index.html, Chapter 8: "Information Superiority and Space," www.dod.mil/execsec/adr2001/Chapter08.pdf (URL inactive)
9. "Obama tells intelligence chiefs to draw up cyber target list—full document text: Eighteen-page presidential memo reveals how Barack Obama has ordered intelligence officials to draw up a list of potential overseas targets for US cyber attacks." *The Guardian*, June 7, 2013, www.guardian.co.uk/world/interactive/2013/jun/07/obama-cyber-directive-full-text
10. This section uses verbatim extracts of articles published by M. E. Kabay originally published in *Network World Security Strategies*. Used with permission.
11. Jesse Berst, "The Electricity Economy: New Opportunities from the Transformation of the Electric Power Sector," White Paper, Global Environment Fund & Global SmartEnergy, August 2008, 55, www.terrawatts.com/electricity-economy.pdf p. 12.
12. Berst, "The Electricity Economy," p. 19.
13. Robert T. Marsh, *Critical Foundations: Protecting America's Infrastructures*, U.S. Government, Washington, DC: President's Commission on Critical Infrastructure Protection, 1997, 192. www.fas.org/sgp/library/pccip.pdf
14. M. E. Kabay, "Attacks on Power Systems: Data Leakage, Espionage, Insider Threats, Sabotage," *Network World Security Strategies*, September 13, 2010, www.mekabay.com/nwss/828c_attacks_on_power_systems_(2).pdf
15. M. E. Kabay, "Attacks on Power Systems: Hackers, Malware," *Network World Security Strategies*, September 8, 2010, www.mekabay.com/nwss/828c_attacks_on_power_systems_(1).pdf
16. ENISA, "Stuxnet Analysis," Press Release, July 10, 2010, www.enisa.europa.eu/media/press-releases/stuxnet-analysis
17. ENISA, "New Report on Smart Grids Cyber Security Measures; A Risk-Based Approach Is Key To Secure Implementation, According to EU Agency ENISA," Press Release, December 19, 2012, www.enisa.europa.eu/media/press-releases/smart-grids-cyber-security-measures-a-risk-based-approach-is-key-to-secure-implementation
18. See www.presidency.ucsb.edu/ws/?pid=24180
19. S. M. Parker, "Information and Finance: A Strategic Target," CommSec, 1997, http://all.net/books/iw/iwarstuff/www.commsec.com/security/infowarfare.htm
20. Federal Reserve Financial Services, "All Fedwire Participants," 2012, https://www.fededirectory.frb.org/fpddir.txt
21. Federal Reserve Board, "Fedwire Funds Sevice 2011 Annual Summary," 2011, www.federalreserve.gov/paymentsystems/fedfunds_ann.htm
22. Clearing House Interbank Payments System, "About CHIPS," www.chips.org/about/pages/033738.php
23. SWIFT "Company Information," www.swift.com/about_swift/company_information/index.page

24. SWIFT, "SWIFT in Figures—FIN Traffic," 2012, www.swift.com/about_swift/company_information/swift_in_figures/archive/2012/SIF_2012_09.pdf (URL inactive)
25. Ann Zimmerman and Miguel Bustillo, "'Flash Robs' Vex Retailers." *Wall Street Journal*, October 21, 2011, http://online.wsj.com/article/SB10001424052970203752604576643422390552158.html
26. M. E. Kabay, "Critical Thinking and Disintermediation," 2007, www.mekabay.com/opinion/critical_thinking.pdf
27. Numbers 13:16, 17
28. Public Law 104-294, "Economic Espionage Act of 1996"; www.law.cornell.edu/usc-cgi/get_external.cgi?type=pubL&target=104-294
29. "Counterintelligence: What Are We Protecting?" Defense Security Service, 1998, www.dss.mil/portal/ShowBinary/BEA%20Repository/new_dss_internet/isp/count_intell/what_protecting.html or http://tinyurl.com/5fwafb
30. T. L. Thomas, "Like Adding Wings to the Tiger: Chinese Information War Theory and Practice," Foreign Military Studies Office, Fort Leavenworth, KS, 2000; www.iwar.org.uk/iwar/resources/china/iw/chinaiw.htm
31. B. Gertz, "Hackers Linked to China Stole Documents from Los Alamos," *Washington Times*, August 3, 2000, p. 1
32. Shen Weiguang, "Checking Information Warfare Epoch Mission of Intellectual Military," *Jiefangjun Bao*, February 2, 1999, p. 6, as translated and downloaded from the Foreign Broadcast Information System (FBIS) Web site on February 17, 1999; www.opensource.gov (registration restricted to U.S. federal, state and local government employees and contractors).
33. Wei Jencheng, "New Form of People's Warfare," *Jiefangjun Bao*, June 11, 1996, p. 6, as translated and reported in FBIS-CHI-96-159, August 16, 1996.
34. Shen Weiguang (1995). "Focus of Contemporary World Military Revolution—Introduction to Research in IW," *Jiefangjun Bao*, November 7, 1995, p. 6, as translated and reported in FBIS-CHI-95-239, December 13, 1995, pp. 22–27.
35. M. E. Kabay, "US DoD Annual Estimates of Information Warfare Capabilities and Commitment of the PRC 2002-2011," 2013, www.mekabay.com/overviews/dod_prc_iw.pdf
36. Qianjin Bao, December 10, 1999, provided by William Belk via email to Timothy L. Thomas. According to Mr. Thomas, Mr. Belk is the head of a skilled U.S. reservist group that studies China.
37. Quotation from S. H. Verstappen, *The Thirty-Six Strategies of Ancient China* (Books and Periodicals, 2000). As described at www.chinastrategies.com/home36.htm
38. Zhang Zhenzhong and Chang Jianguo, "Train Talented People at Different Levels for Information Warfare," *Jiefangjun Bao*, February 2, 1999, as translated and downloaded from FBIS Website on February 10, 1999.
39. A. D. Sofaer et al., "A Proposal for an International Convention on Cyber Crime and Terrorism," 2000, www.iwar.org.uk/law/resources/cybercrime/stanford/cisac-draft.htm
40. G. W. Bush, Executive Order Establishing Office of Homeland Security, 2001, http://georgewbush-whitehouse.archives.gov/news/releases/2001/10/20011008-2.html

41. U.S. Department of Homeland Security, "The Secretary's Five Goals," 2008, www.dhs.gov/xabout/gc_1207339653379.shtm (URL inactive)
42. M. Chertoff, "Remarks by Secretary Michael Chertoff and President of the Supreme Court of Israel Dorit Beinisch to the Heritage Foundation's Civil Rights and the War on Terror: Dilemmas and Challenges Event," April 30, 2008, www.dhs.gov/xnews/speeches/sp_1209741455799.shtm (URL inactive)
43. "FBI: Cyber-Terrorism a Real and Growing Threat to U.S." *Homeland Security News Wire*, March 5, 2010, www.homelandsecuritynewswire.com/fbi-cyber-terrorism-real-and-growing-threat-us
44. Lucian Constantin, "Report: Flame Part of US-Israeli Cyberattack Campaign against Iran." *Network World*, June 20, 2012, www.networkworld.com/news/2012/062012-report-flame-part-of-us-israeli-260353.html
45. M. Curtin and J. Dolske, "A Brute-Force Search of DES Keyspace," 1998, www.interhack.net/pubs/des-key-crack; "Cracking DES: Secrets of Encryption Research, Wiretap Politics & Chip Design—How Federal Agencies Subvert Privacy: Frequently Asked Questions (FAQ) About the Electronic Frontier Foundation's 'DES Cracker' Machine," Electronic Frontier Foundation, 1998, http://w2.eff.org/Privacy/Crypto/Crypto_misc/DESCracker/19980716_eff_des.faq or http://tinyurl.com/68thws; and "RSA Code-Breaking Contest Again Won by Distributed.Net and Electronic Frontier Foundation (EFF): DES Challenge III Broken in Record 22 Hours," Electronic Frontier Foundation, 1999, http://w2.eff.org/Privacy/Crypto/Crypto_misc/DESCracker/HTML/19990119_deschallenge3.html or http://tinyurl.com/5n3gqf
46. E. Rouse, "Psychological Operations/Warfare," date unknown; www.psywarrior.com/psyhist.html
47. K. Poulson, "The Backhoe: A Real Cyberthreat," *Wired*, January 19, 2006, www.wired.com/science/discoveries/news/2006/01/70040; also CGA "CGA DIRT Analysis and Recommendations for Calendar Year 2005," Common Ground Alliance Damage Information Reporting Tool, 2005, www.commongroundalliance.com/TemplateRedirect.cfm?Template=/ContentManagement/ContentDisplay.cfm&ContentFileID=3269 or http://tinyurl.com/43obmo
48. K. Kratovac, "Ship's Anchor Caused Cut in Internet Cable: Unusual Cuts Led to Disruptions in Services, Slowed Down Businesses," MSNBC Technology and Science/Internet, February 8, 2008, www.msnbc.msn.com/id/23068571/
49. United States Strategic Command, "U. S. Cyber Command," Fact Sheet, 2012, www.stratcom.mil/factsheets/Cyber_Command/
50. Scott McNabb, "AFCYBER Takes Part in Second USCYBERCOM Cyber Flag Exercise." U.S. Air Force Space Command press release, November 21, 2012, updated November 29, 2012, www.afspc.af.mil/news/story.asp?id=123327388
51. E. Bumiller, "Pentagon Expanding Cybersecurity Force to Protect Networks Against Attacks," *The New York Times*, January 27, 2013, www.nytimes.com/2013/01/28/us/pentagon-to-beef-up-cybersecurity-force-to-counter-attacks.html
52. NATO, "Collaborating against Cyber Threats," in the Web page "Defending against Cyber Attacks," NATO Website, June 4, 2013, www.nato.int/cps/en/natolive/75747.htm
53. NATO, "NATO and Cyber Defence," NATO Website, www.nato.int/cps/en/SID-BB17E53B-F3456C51/natolive/topics_78170.htm

CHAPTER 15

PENETRATING COMPUTER SYSTEMS AND NETWORKS

Chey Cobb, Stephen Cobb, M. E. Kabay, and Tim Crothers

15.1 MULTIPLE FACTORS INVOLVED IN SYSTEM PENETRATION 15·1

15.1.1 System Security: More than a Technical Issue 15·2
15.1.2 Organizational Culture 15·2
15.1.3 Chapter Organization 15·3

15.2 NONTECHNICAL PENETRATION TECHNIQUES 15·3

15.2.1 Misrepresentation (Social Engineering) 15·3
15.2.2 Incremental Information Leveraging 15·6

15.3 TECHNICAL PENETRATION TECHNIQUES 15·7

15.3.1 Data Leakage: A Fundamental Problem 15·7
15.3.2 Intercepting Communications 15·8
15.3.3 Breaching Access Controls 15·15
15.3.4 Spying 15·20
15.3.5 Penetration Testing, Toolkits, and Techniques 15·20
15.3.6 Penetration via Websites 15·27
15.3.7 Role of Malware and Botnets 15·31
15.3.8 Sophisticated Attackers 15·32

15.4 POLITICAL AND LEGAL ISSUES 15·36

15.4.1 Exchange of System Penetration Information 15·36
15.4.2 Full Disclosure 15·36
15.4.3 Sources 15·38
15.4.4 Future of Penetration 15·39

15.5 SUMMARY 15·40

15.6 FURTHER READING 15·41

15.7 NOTES 15·42

15.1 MULTIPLE FACTORS INVOLVED IN SYSTEM PENETRATION. Although penetrating computer systems and networks may sound like a technical challenge, most information security professionals are aware that systems security has both technical and nontechnical aspects. Both aspects come into play when people attempt to penetrate systems. Both aspects are addressed in this chapter, which is not an instruction guide on how to penetrate systems, but rather a review of the methods and means by which systems penetrations are accomplished.

15.1.1 System Security: More than a Technical Issue. The primary nontechnical factor in system security and resistance to system penetration is human behavior, which can defeat just about any technical security measure. More than anything else, security depends on human beings to understand and carry out security procedures. Consequently, information system (IS) security must be integral to the culture of any organization employing an information system. Without security, systems and networks will not be able resist attempts at penetration.

Often security is represented as a structure of concentric circles. Protection of the central, secured element is then dependent on the barriers imposed by each successive ring. These barriers can be physical or figurative, but the goal of IS security is to protect the integrity, confidentiality, and availability of information processed by the system. This goal is reached using identification, authentication, and authorization. *Identification* is a prerequisite, with each user required to proffer an identifier (ID) that is included in the authorization lists of the system to be accessed. *Authentication* consists of proving that the user really is the person to whom the ID has been assigned. *Authorization* consists of defining what a specific user ID, running specified programs, can legally do on the system. The security perimeter can be penetrated by compromising any of these functions. Chapters 28 and 29 in this *Handbook* discuss identification and authentication in detail.

The trend toward distributed and mobile computers, often utilizing the global networking capability of the Internet, makes it hard to know where to draw these concentric circles of protection. Indeed, the barriers to penetration need to be extended along lines of communication, encompassing end points of the network, which may be geographically dispersed.

15.1.2 Organizational Culture. An organization's general attitude toward security is the key to an effective defense against attack. Security is difficult to sell, especially to an organization that has never experienced a significant problem. (Ironically, the better the defenses, the less evidence there is of their utility.) A basic principle of security is that practitioners must act as if they are paranoid, continuously on guard against attacks from any direction. Many organizations view security precautions as an attack on the integrity of employees. Wearing badges, for example, sometimes is viewed as dehumanizing and offensive. This attitude leads to absurdities, such as having only visitors wear badges. If only visitors wear badges, then taking off the badge automatically reduces the likelihood that a dishonest intruder will be challenged.

Some individual employees also consider security precautions as personally offensive. For example, locking a terminal or workstation when leaving it for a few minutes may be seen as evidence of distrust of other employees. Refusing to allow piggybacking—that is, permitting several colleagues to enter a restricted area on one access card—may be seen as insufferably rude. Where employees are taught to be open and collegial, securing removable computer media and paperwork at night can seem insulting.

These conflicts occur because years of socialization, starting in infancy, are diametrically opposed to the tenets of information security. Politeness in a social context is a disaster in a secure area; for instance, piggybacking into a computer room impairs the accuracy of audit trails kept by the access-control computers. Lending someone a car is kind and generous, but lending someone a user ID and a personal password is a gross violation of responsibility. Chapters 49 and 50 in this *Handbook* discuss psychological aspects of changing corporate culture to support information security.

Carrying out effective security policies and procedures must resolve these conflicts between normal standards of politeness and the standards required in a secure environment. Organizations must foster open discussion of the appropriateness of security procedures, so that employees can voluntarily create a corporate culture conducive to protection of corporate information. Chapters 44, 45, 48, and 51 in this *Handbook* specifically discuss policy issues.

Beyond this, organizations need to be aware of the security posture and attitudes of those with whom they network. These days it is quite possible for one organization's system to be operated, or even owned, by another. And people from many different organizations may be using the same network. A culture of security must permeate all of the organizations that have access to a system, otherwise points of weakness will exist, thus increasing the probability that attempts to penetrate the system will succeed.

15.1.3 Chapter Organization. Section 15.2 looks at methods of tricking people into allowing unauthorized access to systems (Chapter 19 in this *Handbook* explores social engineering in more depth). Section 15.3 examines technical measures for overcoming security barriers and specific techniques (*exploits*) for penetration, while Section 15.4 describes legal and political aspects of system penetration.

15.2 NONTECHNICAL PENETRATION TECHNIQUES. Although the penetration of information systems is often portrayed as the work of the technically adept, many successful penetrations have relied on human factors, such as gullibility and venality. Both are exploited by would-be system penetrators.

15.2.1 Misrepresentation (Social Engineering). Social engineering relies on falsehood. Lies, bribes, and seduction can trick honest or marginally dishonest employees into facilitating a penetration. An attacker might trick an employee into revealing login and authentication codes or even into granting physical access to an otherwise secure site. System penetration can then be accomplished by numerous means, from walking up to an unsecured workstation, to installing Trojan code or a network packet-sniffing device. (Both of these technologies are discussed in more detail later in this chapter.)

15.2.1.1 Lying. Telling lies is a technique often used by persons intent on obtaining unauthorized access to a system. One can obtain valuable information about a system and its defenses by telling lies. Many lies work by playing on the natural human tendency to interpret the world by our internal model of what is most likely. Social psychologists call this model the *schema.* Well-dressed businesspeople who walk briskly and talk assertively are probably what they seem. In a phone conversation, a person who sounds exasperated, impatient, and rude when demanding a new password is probably an exasperated, impatient, and rude employee who has forgotten a password. Unfortunately, many criminals know, sometimes instinctively, how to exploit these interpretations to help get them into secured systems.

Another technique, often used in concert with lying, is to escape notice and avoid suspicion by simply blending in. The way we perceive, or fail to perceive, details are referred to by social psychologists as the figure-ground problem. The normal becomes the background, and the objects of our attention become figures standing out from the ground. The schema influences what is noticed; only deviations from expectation spark figure-ground discrimination. Criminal hackers take advantage of this effect by fading into the background while penetrating security perimeters.

15.2.1.2 Impersonating Authorized Personnel. Criminal hackers and unscrupulous employees call security personnel, operators, programmers, and administrators to request user IDs, privileges, and even passwords. (This is one reason that the telephone is a poor medium for granting security privileges; if staff members were trained to refuse requests made over the phone, many attempts to penetrate systems could be thwarted.) In sites where employees wear ID badges, intruders have a hard time penetrating physical security by posing as employees. However, physical security in these cases depends on the cooperation of all authorized personnel to challenge everyone who fails to wear a badge. This policy is critically important at entry points. To penetrate such sites physically, criminals must steal or forge badges or work with confederates to obtain real but unauthorized badges.

Sites where physical security includes physical tokens, such as cards for electronic access control, are harder for criminals to penetrate. They must obtain a real token, perhaps by theft or by collusion with an employee. Perimeter security depends on keeping the access codes up to date so that cards belonging to ex-employees are inactivated. Security staff must immediately inactivate all cards reported lost. In addition, it is essential that employees not permit *piggybacking,* the act of allowing another person, possibly unauthorized, to enter a restricted zone along with an authorized person. Too often, an employee, in an act of politeness, will permit others to enter a normally locked door as he or she exits. Once inside a building, criminals can steal valuable information that will allow later penetration of the computer systems from remote locations. This is often accomplished by impersonating third-party personnel.

Even if employees are willing to challenge visitors in business suits, it may not occur to them to interfere with people who look as if they are employees of an authorized support firm. For example, thieves often have succeeded in entering a secured zone by dressing like third-party computer technicians or office cleaners. Few employees will think of checking the credentials of a weary technician wearing grimy overalls, an authentic-looking company badge, a colorful ID card, and a tool belt. When a suitable-looking individual claims to have been called to run diagnostics on a workstation, many nontechnical employees will acquiesce at once, seizing the opportunity to grab a cup of coffee or to chat with colleagues. Minutes later, the thief may have copied sensitive files or installed a sniffing device (e.g., a keystroke recorder or a network packet sniffer). In one case known to one of the authors (MK), a criminal was given a workspace and a network connection in a large bank and allowed to work unmolested and unchallenged for several months on a "secret project." It was only when an alert security guard realized that no on one in the office knew who this person was that she challenged the intruder and broke the scam.

15.2.1.3 Intimidation. A technique related to impersonation of authorized or third-party personnel is intimidation. Someone claiming to be a person in a position of authority displays irritation or anger at delays in granting an unauthorized deviation from policy, such as communicating a password over the phone to a person of unauthenticated identity. The attackers indirectly or directly threaten alarming consequences (e.g., delays of critical repairs, financial losses, disciplinary actions) unless they are granted restricted information or access to secured equipment or facilities.

15.2.1.4 Subversion. People make moral choices constantly. There is always a conscious or unconscious balancing of alternatives. Criminal hackers try to reach

their goals by changing the rules so that dishonesty becomes more acceptable to the victim than honesty.

15.2.1.5 Bribery. A lot of industrial and commercial information has a black market value. The same is true of personally identifiable information that can be used to commit fraud and identity theft. The price of a competitor's engineering plans or customer database may be a year's salary for a computer operator responsible for making backups. There is little likelihood that anyone would notice the subverted operator copying a backup at 3:00 A.M. or a secretary taking an extra compact disc out of the office. Many organizations have failed to install software to prevent a manager sending electronic mail with confidential files to a future employer.

That industrial espionage, with or without state sponsorship, is a thriving business is a fact that is now widely—and sometimes quite openly—acknowledged.[1] Building a corporate environment in which employees legitimately feel themselves to be part of a community is a bulwark against espionage. When respect and a sense of exchange for mutual benefits inform the corporate culture, employees will rebuff spies or even entrap them, but the disgruntled employee whose needs are not addressed is a potential enemy.

15.2.1.6 Seduction. Sometimes criminal hackers and spies have obtained confidential information, including access codes, by tricking employees into believing that they are loved. This lie works well enough to allow access to personal effects, sometimes after false passion or drugs have driven the victim into insensibility. It is not unknown for prostitutes to seduce men from organizations that they and their confederates are seeking to crack. Rifling through customers' wallets can often uncover telltale slips bearing user IDs and passwords.

No one can prevent all such abuse. People who are enthralled by expert manipulators will rarely suspect that they are being used as a wedge through a security perimeter. Along with a general increase in security consciousness, staff members with sensitive codes must become aware of these techniques so that they may be less vulnerable. Perhaps then they will automatically reject a request for confidential information or access codes.

15.2.1.7 Extortion. Criminals can threaten harm if their demands are not met. Threaten someone's family or hold a gun to their head and few will, or should, resist a demand for entry to a secured facility or for a login sequence into a network. Some physical access-control systems include a duress signal that can be used to trigger a silent alarm at the monitoring stations. The duress signal requires a predetermined, deliberate action on the part of the person being coerced into admitting unauthorized personnel. This action may be adding an extra number to the normal pass code, pressing the pound sign (#) twice after entering the code, or entering 4357 (H-E-L-P) into the keypad. The duress signal covertly notifies security that an employee is being forced to do something unwillingly. Security can then take appropriate action.

15.2.1.8 Blackmail. Blackmail is extortion based on the threat of revealing secrets. An employee may be entrapped into revealing confidential data, for example, using techniques just described. Classic blackmail includes seduction followed by pictures in flagrante delicto, which the criminals then threaten to reveal. Sometimes a person can be framed by fabricated evidence; a plausible but rigged image of venality can ruin a career as easily as truth. Healthy respect for individuals and social bonds

among employees, supervisors, and management can make it difficult for blackmailers to succeed. If employees who are victims of a blackmail attempt feel they can inform management without suffering inappropriately negative consequences, the threat may be mitigated to a certain degree. Perhaps the last, best defense against blackmail is honesty. The exceptionally honest person will reject opportunities that lead to victimization through blackmail and will laugh at fabrications, trusting friends and colleagues to recognize lies when they hear them.

15.2.1.9 Insiders. Many of the world's largest and most daring robberies have, upon examination, turned out to be inside jobs. The same is true of system penetrations. Although many of the just-described techniques can be used to obtain help from the inside, some are made possible by people on the inside who decide, for whatever reason, to aid and abet criminal hackers. For example, a dishonest employee may actively seek to sell access for personal gain. Organizations should try to be alert to this eventuality, but there is very little defense against thoroughly dishonest employees when the only overt act needed to open the gates from the inside is to pass system credentials to an outsider. Chapter 13 in this *Handbook* specifically addresses insider crime.

15.2.1.10 Human Target Range. Organizations should not underestimate the range of targets at which the described techniques may be directed. Although the terms *employees*, *authorized personnel*, and *third-party personnel* are used in the preceding paragraphs, the target range includes all manner of vendors, suppliers, and contractors as well as all levels of employees—from software and hardware vendors, through contract programmers, to soft-drink vendors and cleaning staff. It may even include clients and customers, some of whom possess detailed knowledge of the organization's operations. Employees at every level are likely to be computer literate, although with varying degrees of skill. For example, it is quite possible that someone working as a janitor today knows how to operate a computer skillfully and may even know how to surf hacking sites on the Web and download penetration tools. Indeed, a janitor may have obtained the job specifically with the intent of engaging in industrial espionage, data theft, or sabotage.

In short, anyone who comes into contact with the organization has the potential to provide an attacker with information useful in the preparation and execution of an attack. The human targets of a social engineering attack may not, on an individual basis, possess or divulge critical information, but each may provide clues—pieces of the puzzle—an aggregation of which can lead to successful penetration and compromise of valuable data and resources. Use of this process is a hallmark of some of the most successful criminal hackers. The term *incremental information leveraging* was coined for this use of less valuable data to obtain more valuable data.[2]

15.2.2 Incremental Information Leveraging. By gathering and shrewdly utilizing small and seemingly insignificant pieces of information, it is possible to gain access to much more valuable information. This technique of incremental information leveraging is a favorite tool of hackers, both criminal and noncriminal. One important benefit of the tool that is particularly appreciated by criminal hackers is the low profile it presents to most forms of detection. By accumulating seemingly innocuous pieces of information over a period of time, and by making intelligent deductions from them, it is possible to penetrate systems to the highest level.

A prime example of this approach is seen in the exploits of Kevin Mitnick, who served almost five years behind bars for breaking into computers, stealing data, and

abusing electronic communication systems. Illegal acts committed by Mitnick include the 1981 penetration of Computer System for Mainframe Operations (COSMOS), a Pacific Bell facility in downtown Los Angeles. COSMOS was a centralized database used by many U.S. phone companies for controlling basic recordkeeping functions. Mitnick and others talked their way past a security guard and located the COSMOS computer room. They stole lists of computer passwords, operating manuals for the COSMOS system, and combinations to the door locks on nine Pacific Bell central offices. Mitnick later employed knowledge of phone systems and phone company operations to penetrate systems at Digital Equipment Corp. (DEC).

Since his release in January 2000, Mitnick has spoken about information security before Congress and at other public venues. He described social engineering as such a powerful tool that he "rarely had to resort to a technical attack."[3] As to technique, he stated, "I used to do a lot of improvising . . . I would try to learn their internal lingo and tidbits of information that only an employee would know." In other words, by building up knowledge of the target, using a lot of information that is neither protected nor proprietary, it is possible to gain access to that which is both proprietary and protected. The power of incremental information leveraging is the equivalent of converting a foot in the door into an invitation to come inside.

Protection against incremental information leveraging and all other aspects of social engineering begins with employee awareness. Employees who maintain a healthy skepticism toward any and all requests for information provide a strong line of defense. Another powerful defense mechanism, highlighted by Mitnick, is the use of telephone recording messages, such as "This message may be monitored or recorded for training purposes and quality assurance." An attacker who hears a message like this may think twice about proceeding with attempts to use voice calls to social engineer information from the target.

15.3 TECHNICAL PENETRATION TECHNIQUES. Technical penetration attacks may build on data obtained from social engineering, or they may be executed on a purely technical basis. Techniques used include eavesdropping, either by listening in on conversations or by trapping data during transmissions, and breaches of access controls (e.g., trying all possible passwords for a user ID or guessing at passwords). Weaknesses in the design and implementation of information systems, such as program bugs and lack of input validation, also may be exploited in technical attacks. Unfortunately, weaknesses of this nature abound in the realm of the Internet, even as more and more organizations increase their Internet connectivity, thus creating more and more potential penetration points.

15.3.1 Data Leakage: A Fundamental Problem. Unfortunately, for information security (INFOSEC) specialists, it is impossible, even in theory, to prevent the unauthorized flow of information from a secured region into an unsecured region. The imperceptible transfer of data without authorization is known as *data leakage*. Technical means alone cannot suppress data leakage.

Consider a tightly secured operating system or security monitor that prevents confidential data from being copied into unsecured files. Workstations are diskless, there are no printers, employees do not take disks into or out of the secured facility, and there are strict restrictions on taking printouts out of the building. These mechanisms should suffice to prevent data leakage.

Not really.

Anyone with a penchant for mnemonics or with a photographic memory could simply remember information and write it down after leaving the facility. And it is extremely difficult to prevent employees from writing notes on paper and concealing them in their clothing or personal possessions when they leave work. Unless employees are strip-searched, no guard can stop people with crib sheets full of confidential data from walking out of the building. Indeed, this is how Vasili Mitrokhin, head archivist of the KGB's First Chief Directorate, perpetrated the largest breach of KGB security ever, by smuggling thousands of handwritten copies of secret documents out of the KGB headquarters in Moscow in his shoes, socks, and other garments.[4]

Another means of data leakage is steganography, hiding valuable information in plain sight among large quantities of unexceptional information. For example, a corrupt employee determined to send a confederate information about a chemical formula could encode text as numerical equivalents and print these values as, say, the fourth and fifth digits of a set of engineering figures. No one is likely to notice that these numbers contained anything special. The more digitally inclined can use steganography software, freely available on the Internet, to hide data in image files.

The unauthorized transfer of information cannot be absolutely prevented because information can be communicated by anything that can fluctuate. Theoretically, one could transfer data to a confederate by changing the position of a window shade (slow but possible). Or one could send ones and zeroes by the direction of oscillation of a tape reel; or one could send coded information by the choice of music. Even if a building were completely sealed, it would still leak heat outward or transfer heat inward—and *that* would be enough to carry information. In practical terms, system managers can best meet the problem of data leakage by a combination of technical protection and effective management strategies.

It is also important to realize that a significant amount of data leakage occurs through innocuous intending communications from employees. It is common for employees to discuss small aspects of their jobs and work information without realizing the implications and ability of others to collate this information into far more substantial amounts. This has become particularly prevalent with the modern advent of social media. Significant amounts of information about an organization and its internal workings can be derived from employee social media pages. Consequently, social media has become a rich source of data for attackers wishing to target an organization. This approach is particularly common among attackers that target an organization through its employees using techniques such as spear-phishing.

Data loss prevention (DLP), discussed briefly in Chapter 13 in this *Handbook*, is now an established set of techniques, with numerous tools available to enforce restrictions on unauthorized data transfers to storage devices and external sites. Nonetheless, vigilance over human behavior and the effective configuration and real-time or at least frequent analysis of log records remain essential components of effective DLP.

15.3.2 Intercepting Communications. Criminal hackers and dishonest or disgruntled employees can glean access codes and other information useful to their system penetration efforts by monitoring communications. These might be between two workstations on a local area or wide area network, between a remote terminal and a host such as a mainframe, or between a client and a server on the Internet. Attackers can exploit various vulnerabilities of communications technologies. The shift to Transmission Control Protocol/Internet Protocol (TCP/IP)–based Internet communications over the last decade has brought many more communications streams into the target

range of would-be penetrators. For basic information about data communications, see Chapter 5 in this *Handbook*.

15.3.2.1 Wiretapping. Wiretapping consists of intercepting the data stream on a communications channel (even if that channel is not wire; e.g., fiber optic cable can also be tapped, as can wireless communications, although the latter sometimes are said to be sniffed rather than tapped).

15.3.2.2 Asynchronous Connections. Point-to-point data connections (e.g., using telephone modems or serial devices) have all but disappeared, but they are relatively easy to tap. Physical connection at any point on twisted pair or multiwire cables allows eavesdropping on conversations; if the line is being used for data communications, a monitor is easily attached to display and record all information passing between a node and its host. Asynchronous lines in large installations often pass through patch panels, where taps may not be noticed by busy support staff, as they manage hundreds of legitimate connections. Such communications usually use phone lines for distances beyond a few hundred meters (or about 1,000 feet).

Wiretappers must use modems configured for the correct communications parameters, including speed, parity, number of data bits, and number of stop bits, but these parameters are easy to find out by trial and error.

Countermeasures include:

- Physical shielding of cables and patch panels
- Multiplexing data streams on the same wires
- Encryption of data flowing between nodes and hosts

15.3.2.3 Synchronous Communications. Because synchronous modems are more complex than asynchronous models and because their bandwidths (maximum transmission speeds) are higher, they are less susceptible to attack, but they are not risk-free.

15.3.2.4 Dial-up Phone Lines. Used for both data and voice communications, dial-up lines supplied by local telephone companies and long-distance carriers are vulnerable to wiretapping. Law enforcement authorities and telephone company employees can install taps at central switching. Criminals can tap phone lines within a building at patch panels, within cabling manifolds running in dropped ceilings, below raised floors, or even in drywall. They also can tap at junction boxes where lines join the telephone company's external cables.

The same countermeasures apply to phone lines as to asynchronous or synchronous data communications cables.

15.3.2.5 Leased Lines. Leased lines use the same technology as dial-up (switched) lines, except that the phone company supplies a fixed sequence of connections rather than random switching from one central station to another. There is nothing inherently more secure about a leased line than a switched line; on the contrary, it is easier to tap a leased line at the central switching station because its path is fixed. However, leased lines usually carry high-volume transmissions. The higher the volume of multiplexed data, the more difficult it is for amateur hackers to disentangle the data streams and make sense of them. At the high end of leased line bandwidth

(e.g., carriers such as T1, T2, etc.), the cost of multiplexing equipment makes interception prohibitively expensive for all but professional or government wiretappers.

Data encryption provides the best defense against wiretapping on leased lines.

15.3.2.6 Long-Distance Transmissions. Dial-up and leased lines carry both short-haul and long-distance transmissions. The latter introduce additional points of vulnerability. Microwave relay towers carry much of the long-distance voice and data communications within a continent. The towers are spaced about 40 kilometers (25 miles) apart; signals spread out noticeably over such distances. Radio receivers at ground level can intercept the signals relayed through a nearby tower, and because microwaves travel in straight lines, rather than following the curvature of the earth, they eventually end up in space, where satellite receivers can collect them. The difficulty for the eavesdropper is that there may be thousands of such signals, including voice and data, at any tower. Sorting out the interesting ones is the challenge. However, given sufficient computing power, such sorting is possible, as is targeting of specific message streams. *Spread-spectrum* transmission, or *frequency hopping*, is an effective countermeasure.

15.3.2.7 Packet-Switching Networks. Packet-switching networks, including historical X.25 carriers such as Telenet, Tymnet, and Datapac, used packet assembler-disassemblers (PADs) to group data into packets addressed from a source to a destination. If data traveled over ordinary phone lines to reach the network, interception could occur anywhere along these segments of the communications link. However, once the data were broken up into packets (whether at the customer side or at the network side), wiretappers had a difficult time making sense of the data stream.

15.3.2.8 Internet Connections. TCP/IP connections are no harder to tap than any others, and they carry an ever-increasing array of data, from e-commerce traffic to television broadcasts and voice communications, the latter using Voice over Internet Protocol (VoIP). (See Chapter 34 in this *Handbook*.) Unless the data stream is encrypted, there are no special impediments to wiretappers. Although the tapping of fiber optic cable requires more specialized equipment than the tapping of copper cables, it is possible.

15.3.2.9 LAN Packet Capture. Local area networks (LANs) are similar to packet-switching networks: both network protocols send information in discrete packages, either over cables or radio waves. Each package has a header containing the address of its sender and of its intended recipient. Packets are transmitted to all nodes on a segment of a LAN. For more information about LANs, see Chapter 25 in this *Handbook*.

Normally a node is restricted to interpreting only those packets that are intended for it alone. However, it is possible to place devices in "promiscuous mode," overriding this restriction. This can be done with software that surreptitiously converts a device, such as an end user workstation, into a listening device, capturing all packets that reach that node. Of course, network administrators can intentionally create a packet-capturing workstation for legitimate purposes, such as diagnosing network bottlenecks. It is also possible to connect specialized hardware called *LAN monitors* to the network, either with or without permission, for legitimate or illegitimate purposes. Sometimes called *network sniffers*, these devices and programs range from basic freeware to expensive commercial packages that can cost tens of thousands of dollars for a network with

hundreds of nodes. (The term *sniffer*, although in common use, is a registered trademark, as *Sniffer*®, of Network General Corporation.)

The more sophisticated packet-sniffing programs allow the user to configure profiles for capture; for example, the operator can select packets passing between a host and a system manager's workstation. Such programs allow an observer to view and record everything seen and done on a workstation, including logins or encryption keys sent to a server.

Packet sniffing poses a serious threat to confidentiality of data transmissions through LANs. Most sniffing programs do not announce their presence on the network. Although it may not be apparent to the casual observer that a workstation is performing sniffing, it is possible, as a countermeasure, to scan the network for sniffing devices. Stealthier packet-sniffing technology is constantly improving, and tight physical security may be the best overall deterrence.

LAN users concerned about confidentiality should use LAN protocols that provide end-to-end encryption of the data stream or third-party products to encrypt sensitive files before they are sent through the LAN. Routers that isolate segments of an LAN or WAN (wide area network) can help limit exposure to the threat of sniffers.

15.3.2.10 Optical Fiber. Although optical fibers were once thought to be secure against interception, new developments quickly abolished that hope. An attacker can strip an optical fiber of its outer casing and bend it into a hairpin with a radius of a few millimeters (1/8 inch); from the bend, enough light leaks out to duplicate the data stream. Bryan Betts, writing in *PCWorld*, quoted

> Thomas Meier, the CEO of Swiss company Infoguard. . . . [who] demonstrated the technique on a fibre carrying a VOIP phone call over Gigabit Ethernet. A section of fibre from inside a junction box was looped into a photodetector called a bend coupler, and the call was recorded and then played back on a laptop.
>
> "People claim optical fibre is harder to tap than copper, but the opposite is true—you don't even have to break the insulation, as you would with copper," Meier said. "You can read through the fibre's cladding with as little as half a dB signal loss."
>
> He claimed that suitable bend couplers can be bought off the shelf—or from eBay—for a few hundred dollars, and connected to the extra fibre that is typically left coiled up in junctions boxes for future splicing needs.
>
> He added that the risk is not imaginary or theoretical—optical taps have been found on police networks in the Netherlands and Germany, and the FBI investigated one discovered on Verizon's network in the United States. Networks used by U.K. and French pharmaceutical companies have also been attacked, probably for industrial espionage, he said.[5]

Luckily, most optical trunk cables carry hundreds or thousands of fibers, making it almost impossible to locate any specific communications channel. (The same is not true of fiber cables used to deliver network connectivity to individual homes and offices.) Equipment for converting optical signals into usable data remains costly, discouraging its use by casual criminal hackers.

15.3.2.11 Wireless Communications. Cable-based communications have the advantage of restricting channel access to at least theoretically visible connections. However, the rapid increase in wireless telecommunications in the last decade of the twentieth century and the first years of the twenty-first has routed increasing

amounts of information through a broadcast medium in which access—even unauthorized access—may be invisible to users and system administrators.

15.3.2.12 Wireless Phones. Also referred to as cordless phones, conventional wireless phones broadcast their signals, and the traffic they carry can be detected from a distance. Older cordless phones were analog and susceptible to eavesdropping from such basic devices as walkie-talkies and baby monitors. Children sometimes walked around their suburban neighborhoods with a handset from such a phone turned on; once they walked far enough away from their home to lose the signal, any new dial tone belonged to a neighbor, who might be puzzled to discover a call to the Antipodes on the next phone bill. Today's wireless phone models typically use a different set of frequencies, such as 2.4 gigahertz (GHz). These phones generally use frequency-hopping spread spectrum (FHSS) technology to make unauthorized use more difficult, and to impede eavesdropping. FHSS means the signals hop from frequency to frequency across the entire 2.4-GHz spectrum, making tapping their signals harder but by no means impossible. Wireless phones should not be used for confidential voice or data traffic unless encryption is enabled and activated. An added danger lies in hanging up a cordless phone during a conversation, since the cordless phone's base continues to transmit until switched off.

15.3.2.13 Cellular (Mobile) Phones. Early analog cellular (mobile) phone systems had an expectation of privacy equivalent to that of shouting a message through a megaphone from a rooftop. Calls on such phones were easily intercepted using scanners purchased from local electronics stores. Although encryption is possible on the newer digital cell phones that are now widely used, the encryption is not always turned on due to the burden it imposes on the cell company switching equipment. Check with the carrier before assuming that cell calls are encrypted. Also, bear in mind that, although digital cell phone calls are harder to intercept than analog ones, there is a thriving black market in devices that make such interception possible.

As a rule, confidential information should never be conveyed through cellular phones without encrypting the line or the messages first.

Phil Zimmermann, creator of Pretty Good Privacy (PGP, later GPG) in the early 1990s, created a new service in 2012 called Silent Circle to provide encrypted telephony and Internet access to its subscribers.[6]

15.3.2.14 Wireless Networks. The increasingly popular means of networking computers to networks known as WiFi presents many opportunities for interception of communications. In this context, "wireless network" usually means a data network using the 802.11 standard, which comes in a variety of flavors, such as 802.11b, 802.11g (often collectively referred to as WiFi, which stands for "Wireless Fidelity" and is actually a brand name owned by the trade group WiFi Alliance). Typically, this is the sort of local area network created by plugging a wireless access point into an Ethernet network and a WiFi card or adapter into each computer. Most notebook computers now come with a built-in WiFi adapter.

These WLANs, or wireless local area networks, are relatively cheap and easy to create since they do not require network cabling. That helps to explain why more than 200 million WiFi devices were sold in 2006 and more than half of all U.S. companies have been using WLANs to some degree or other since 2002. But cheap WLANs come with hidden costs, namely security. Every WLAN operates, by its very nature, in loose-lips mode. In a sense, the ease of use comes with ease of abuse. They all broadcast

their traffic into the air, whence it may be overheard by someone with the right set of ears, a legitimate user or a criminal hacker, someone looking for free bandwidth or a war driver. War driving, the practice of driving around town to find wireless access points, is a hobby to some people, and probably not illegal unless done with malicious intent—many notebook computers try to find wireless access points whenever turned on—but bear in mind that laws vary from one country to another. (Those contemplating war driving should check the legal status in their jurisdictions.)

To be a war driver, all that is needed is an old laptop, the right WiFi card, some free software (NetStumbler, e.g.), and an empty Pringles® potato-chip can wired to the WiFi card as an external antenna to boost reception. (The Pringles® can is optional, as is a global positioning system device to mark the location of WiFi access points.) If you drive around with this equipment activated, you will doubtless discover numerous access points in both residential and business districts. If the names of the access points are things like "Linksys" or "netgear," this is an indication that the owner of the network has not changed the default service set identifier (SSID), which tends to be the brand of the wireless access point, broadcast for the world to see, unless the network owner turns off this feature. The name could also be a person, or place, or company, which helps war drivers figure out whose network they are picking up. (One of the authors detected SCHS near the offices of Sample County Health Services.) A program like NetStumbler will also tell you whether the network is using encryption. There has been a steady rise in the percentage of wireless networks using encryption, but it is far less than 10 percent. According to a survey conducted by AT&T in late 2007, one in six small businesses in America that use wireless technology has taken no precautions against wireless threats, and one-third of small businesses indicated that they were unconcerned about wireless data security.[7]

There is more about wireless network security in Chapter 33 in this *Handbook,* but the point is the relative ease with which networks can be tapped. This means WiFi, whether at home, in the office, or at a hot spot, represents a significant category of data leakage and thus a major avenue for systems penetration.

15.3.2.15 Van Eck Phreaking. This attack is named for Wim Van Eck, a Dutch electronics researcher who in 1985 proved to a number of banks that it was possible to read information from their cathode ray tubes (CRTs) at distances of almost a mile away, using relatively simple technology.[8] Because many types of electronic equipment emit radio-frequency signals, receivers that capture these signals can be used to reconstruct keystrokes, video displays, and print streams. Using simple, inexpensive wide-band receivers, criminals can detect and use such emissions at distances of tens or hundreds of meters (yards).

Since radio-frequency signals leak easily through single-pane windows, PCs should never be placed in full view of ground-floor windows (and certainly not facing the windows!). Attenuators that tap the window at irregular intervals can be installed to defeat such leakage. A special double-pane window with inert gas between the panes also can lessen the amount of signals leakage.

Other countermeasures include special cladding of hardware, such as computers and printers, to attenuate broadcast signals. This protection often is referred to by the name of the classified government standard for protection of sensitive military systems, TEMPEST. Although TEMPEST was allegedly a classified code word to begin with, it is now sometimes expanded as *Transient ElectroMagnetic Pulse Emission Standard* or *Telecommunications Electronics Material Protected from Emanating Spurious Transmissions*.

TEMPEST-certified equipment costs many times more than the same equipment without TEMPEST cladding.[9] A less expensive alternative is to use a special device that emits electromagnetic noise that masks meaningful signals. Yet another approach to protection against this threat is to locate systems within buildings, or rooms within buildings, that have been constructed to TEMPEST standards. There are federal regulations concerning the methods of building sensitive compartmented information facilities (SCIFs), and the testing to obtain a TEMPEST rating is quite stringent. These measures include such things as cladding of all walls and ceilings, cladding of all electrical and network cabling, lead-lined doors, and the absence of any external windows.

15.3.2.16 Trapping Login Information. Criminals can capture identification and authentication codes by inserting Trojan horse programs into the login process on a server host and by using macro facilities to record keystrokes on a client node. More commonly today, however, specifically written malware is used to capture keystrokes and transmit them at intervals back to remote systems.

15.3.2.17 Host-Based Login Trojans. A Trojan horse is a program that looks useful but contains unauthorized, undocumented code for unauthorized functions. The name comes from Greek mythology, in which Odysseus (Ulysses in Latin), weary of the never-ending siege of Troy, sailed his ships out of sight as if he and his warriors were giving up, but left a giant wooden horse at the city gates. Entranced by this magnificent peace offering, the Trojans dragged the great horse into the city. During the Trojans' wild celebrations that night, the soldiers Odysseus had secreted in the belly of the hollow horse let themselves out and opened the gates to their army. The Greeks slaughtered all the inhabitants of the city, and the Trojan War was over.

In February 1994, the Computer Emergency Response Team Coordination Center (CERT-CC) at Carnegie Mellon University in Pittsburgh issued a warning that criminal hackers had inserted Trojan horse login programs in hundreds of UNIX systems on the Internet. The Trojan captured the first 128 bytes of every login and wrote them to a log file that was later read by the criminals. This trick compromised about 10,000 login IDs.

Trojan code might be installed on a computer or terminal used by several people (e.g., on a mainframe terminal in the 1970s or in an Internet café today) so that when someone enters a user ID and password to logon, the system—controlled by the Trojan—displays a message such as "Invalid password, try again," and the user does so. This time the login is accepted. The victim continues working, unaware that there is anything unusual going on. The Trojan, installed earlier, simulated the normal login procedure, displaying a semblance of the expected screen and dialog. Once the victim entered a password, the Trojan writes the authentication data to a file and then shows a misleading error message. The spoof program then terminates and the regular program is ready for login.

Such a case occurred in April 1993 in a suburb of Hartford, Connecticut. Shoppers noticed a new automated teller machine (ATM) in their mall. At first, the device seemed to work correctly, disbursing a few hundred dollars to bank card users on demand. It quickly changed to a more sinister mode. Users would insert their bank cards and respond as usual to the demand for their personal identification numbers (PIN). At this point, the new ATM would flash a message showing a malfunction and suggesting that the user try an adjacent bank machine. Most people thought nothing of it, but eventually someone realized that the ATM was not posting the usual "Out of Order"

indicator after these supposed errors. In addition, banks began receiving complaints of a rash of bank card frauds in the immediate area. Investigators discovered that the ATM had no connection to any known bank—that it had been purchased used, along with equipment for manufacturing bank cards. The ATM was a spoof; it was merely collecting the user ID and PIN of every victim for later pickup by the criminals who had installed it without permission in the mall. The criminals were caught after having stolen about $100,000 over a four-week period using fraudulent bank cards.

Chapter 20 in this *Handbook* includes more details about Trojans and other low-tech attacks.

15.3.2.18 Keystroke Logging. Another threat to identification and authentication codes is the ability to record keystrokes for later playback or editing. Most word processing programs provide macro facilities, so named because of their ability to store and output multiple keystrokes, such as the typing of boilerplate text, with one keystroke. More sophisticated terminate-and-stay-resident (TSR) programs can record sequences of commands, and are sometimes used to demonstrate software or to automate quality assurance tests. This technology can also be used to lay in wait on a workstation and record everything the user does with the mouse and types with the keyboard; such programs are sometimes called *keystroke loggers*. Later, the criminal can harvest the records and pick out the login codes and other valuable information.

There are also hardware implementations of keystroke logging. One is a small device inserted between the keyboard and the computer, capturing what is typed and holding it in nonvolatile memory until it can be retrieved. At the time of this writing in May 2013, such devices were widely available through Internet sales for about US$40.

By far the most common form of keystroke logging is done by *hooking* the operating system mechanisms for providing keyboard functionality. This is done in software. There are numerous specific technical techniques to accomplish the hooking process depending on the operating system and hardware. Sometimes the software used to accomplish the keystroke logging is stand-alone, but often it is part of a larger piece of software that also provides C2 (command and control) capabilities. It is also quite common for the keystroke logging functionality to be part of a rootkit in order to avoid detection.

It is possible to defeat the attempted reuse and abuse of login credentials captured by any of these methods by switching to one-time passwords generated by microprocessors. One-time passwords are discussed in Chapter 28 of this *Handbook*. However, both key loggers and Trojans can be deployed to gain unauthorized access to data without resorting to the reuse of passwords.

15.3.3 Breaching Access Controls. Criminals and spies use two broad categories of technical attacks to deduce access phone numbers, user IDs, and passwords: brute-force attacks and intelligent guesswork. In addition, there are ways to manipulate people into revealing their access codes; these techniques are discussed in the section on social engineering.

15.3.3.1 Brute-Force Attacks. Brute-force attacks consist of using powerful computers to try all possible codes to locate the correct ones. Brute force is applied to locating modems, network access points, vulnerable Internet servers, user IDs, and passwords.

15.3.3.2 Demon (War) Dialing. Despite the wholesale shift of data communications to TCP/IP networks and the Internet, modems on dial-up phone lines remain a common, and sometimes forgotten, means of external access to a system. The telephone numbers of any modems connected to hosts or servers or intelligent network peripherals, such as high-end laser printers, are sensitive and should not be posted or broadcast.

Demon dialers are programs that can try every phone number in a numerical range and record whether there is a voice response, a fax line, a modem carrier, or no answer. When phones ring all over an office in numerical order, one at a time, and when there is no one on the line if a phone is picked up, it is undoubtedly the work of someone using a demon dialer. Of course, good demon dialing software accesses the numbers in the target range nonsequentially.

During the heyday of fax machines, some youngsters were reported to have "farmed" entire telephone exchanges during the night, then to have sold the fax numbers for $1 dollar per number to unscrupulous junk-fax services that sold advertisers access to them.

15.3.3.3 Exhaustive Search. The same approach as demon dialing can find user IDs and passwords after a connection has been made. The attacker uses a program that cycles systematically through all possible user IDs and passwords and records successful attempts. The time required for this attack depends on two factors:

1. The keyspace for the login codes
2. The maximum allowable speed for trying logins

In today's technical environment, any inexpensive computer can generate login codes far faster than hosts permit login attempts. Processor speed is no longer a rate-limiting factor. Note that this type of attempt to "guess" passwords is different from password cracking, described elsewhere, which operates on captured or stolen copies of encrypted password files.

15.3.3.4 Keyspace. As discussed in Chapter 7 in this *Handbook*, the keyspace for a code is the maximum number of possible strings that meet the rules of the login restrictions. For example, if user passwords consist of exactly six uppercase or lowercase letters or numbers and the passwords are case-sensitive (i.e., uppercase letters are distinguished from lowercase letters), the total number of possible combinations for such passwords is calculated in this way:

- There are 10 digits and 52 upper- or lowercase letters (in the English alphabet) = 62 possible codes for any of six positions.
- If there are no restrictions on repetition, a string of n characters to be taken from a list of r possibilities for each position will generate r^n possible combinations.
- Thus, in our example, there are 62^6 possible sequences of 62 codes taken in groups of six = 56,800,235,584 (more than 56 billion) possible login codes.

If there are restrictions, the keyspace will be reduced accordingly. For example, if the first character of a password of length six must be an uppercase letter instead of being any letter or number, there are only 26 possibilities for that position instead 62,

thus reducing the total keyspace to $26 \times 62^5 = 23{,}819{,}453{,}632$ (more than 23 billion) possibilities.

15.3.3.5 Rainbow Tables. Cryptanalysts (including criminal cryptanalysts) can generate all possible one-way hashes for the keyspace of any given password rule; brute-force cracking using the tables can thus be accelerated. There are even Websites distributing such tables freely.[10]

15.3.3.6 Login Speed. Generating login codes is not hard. The greatest barrier to brute-force login attacks is interruption in the login whenever the host detects an error. Most operating systems and security monitors allow the administrator to define two types of login delays following errors:

1. A usually brief delay after each failed attempt to enter a correct password
2. A usually long delay after several failed login attempts

Suppose each wrong password entered causes a 1/10th-second delay before the next password can be entered; then for our example involving six repeatable uppercase or lowercase letters or numbers, it would take 5,680,023,558 seconds = 1,577,784 hours $\approx$ 180 years to try every possibility.

Suppose, in addition, that after every fifth failed login attempt, the system were to inactivate the user ID or the modem port for three minutes. Such interference would stretch the theoretical time for a brute-force exhaustive attack to around 650 years.

Should the security manager completely inactivate the ID if it is under attack? If the ID is inactivated until the user calls in for help, user IDs become vulnerable to inactivation by malicious hackers. Attackers need merely provide a bad password several times in a row and the unsuspecting *legitimate* user will be locked out of the system until further notice. A widespread attack on multiple user IDs could make the system unavailable to most users. Such a result would be a denial-of-service attack (see Chapter 18 in this *Handbook*).

Should the port be inactivated? If there are only a few ports, shutting them down will make the system unavailable to legitimate users. This drastic response may be inappropriate—indeed, it may satisfy the intentions of criminal hackers. A short delay, perhaps a few minutes, would likely be sufficient to discourage brute-force attacks.

In all of these examples, the illustrations have been based on exhaustive attacks (i.e., trying every possibility). However, if passwords or other codes are chosen randomly, the valid codes will be uniformly distributed throughout the keyspace. On average, then, according to a principle of statistics called the *Central Limit Theorem*, brute-force searches will have to search half the keyspace. For large keyspaces, the difference between a very long time and half of a very long time will be negligible in practice (e.g., 325 years is not significantly different from 650 years if everyone interested will be dead before the code is cracked).

15.3.3.7 Scavenging Random Access Memory. Not all attacks come from outside agents. Criminals with physical access to workstations, malicious software, or authorized users who can use privileged utilities to read main memory, can scavenge memory areas for confidential information such as login IDs and passwords.

On a workstation using a terminal emulator to work with a host, ending a session does not necessarily unload the emulator. Many emulators have a configurable screen

display buffer, sometimes thousands of lines long. After an authorized user logs off and leaves a terminal, a scavenger can read back many pages of activity, sometimes including confidential information or even login codes. Passwords, however, usually are invisible and therefore not at risk.

If a workstation is part of a client/server system, an application program controlling access may leave residues in random access memory (RAM). A RAM editor, easily available as part of utility packages, can capture and decode such areas as file buffers or input/output (I/O) buffers for communications ports. However, rebooting the workstation after communication is over prevents RAM scavenging by reinitializing memory.

15.3.3.8 Scavenging Cache Files. The same principle can be applied to the various cache and swap files created by the operating system. Cache files are used to keep frequently used data readily available to applications. Swap files store data and code that is moved out of memory onto disk when memory is full. Some operating systems also create hibernation files, writing memory to disk just prior to powering down, and thus enabling a quick resumption of work when the system is powered up. Operating systems may also provide auto-saved recovery files that allow restoration of data after a system error. All of these can be mined for system credentials as well as other valuable data.

15.3.3.9 Scavenging Web History Files. A more recent variation on this scavenging approach is to examine files created by Web browsers. These sometimes contain not only the pages viewed by a user but also the credentials entered to access those pages.

15.3.3.10 Recovering Stored Passwords. Many applications store user passwords locally. Most applications attempt to secure these passwords, unfortunately they are seldom stored with proper encryption methods. It has become common for malicious software to recover these stored passwords and transmit them to remote servers for collection by attackers.

15.3.3.11 Intelligent Guesswork. Users rarely choose random passwords. Much more frequently, passwords are chosen from a subset of all possible strings. Instead of blindly batting at all possible sequences in a keyspace, an attacker can try to reduce the effective keyspace by guessing at more likely selections. Likely selections include canonical passwords, bad passwords, and words selected from a dictionary.

Hardware and software often come from the factory, or out of the box, with user IDs and passwords that are the same for all systems and users.

For example, wireless access points and routers have default user IDs when they ship from the factory. Naturally, these user IDs are set up with the same password. (For example, “admin” on Linksys wireless router devices.) Such systems always include instructions to change the passwords, but, too often, administrators and users neglect to do so. Criminals are familiar with factory presets—most of which are readily discoverable via Google—and exploit them to penetrate systems. The simple routine of changing all canonical passwords prevents hackers from gaining easy access to systems and software.

15.3.3.12 Stealing. Any element of a computer system has potential value to a thief. Therefore every element including hardware, software, media, files, documentation, and printouts must be safeguarded.

15.3.3.13 Data Scavenging. Criminal hackers have few scruples about using other people's property when they enter computer systems; they have none at all concerning using other people's trash. The term *data scavenging* describes the process that acquires information from throw-away sources. Perhaps the most widely known is Dumpster® diving, sorting through whatever an organization discards. Hard-copy printouts, CD-ROMs, tapes, and other assorted data-bearing media often end up in trash containers where they are easily accessible to Dumpster® divers after hours. In some areas, if one visits an office or industrial park at night, one can see half a dozen people rummaging about, sometimes headfirst in Dumpsters®. Criminal hackers use the information thoughtlessly discarded by naïve office workers as a source of procedures, vocabulary, and proper names that can help them impersonate employees over the phone or even in person. The classic example is a discarded internal phone directory, which can provide a social engineer with valuable data to use when making calls to employees. An employee who hesitates to comply with an attacker's bogus request for information over the phone may well be persuaded if the attacker says something like "I understand your hesitation; if it makes you feel more comfortable you can call me back at extension 2645." If 2645 is a legitimate internal extension, the caller gains considerable credence. Of course, the properly trained employee will hang up and make the call to 2645 rather than take the easy option and say, "I guess that's okay then, here is the information you wanted."

Some printouts contain confidential information that can lead to extortion or system penetration. For example, a thief who steals a list of personally identifiable information about patients with HIV infection could torment the victims and extort money. Every piece of paper, or other media to be discarded, should be evaluated for confidentiality. Unless the information is worthless to everyone, employees should shred paper before disposal or arrange to send paper to a bonded service for destruction. The same applies to CD-ROMs and other media.

15.3.3.14 Discarded Magnetic and Optical Media. Discarded paper poses a threat; discarded magnetic and optical media are a disaster. Many organizations fail to teach employees that the normal commands used to delete files do not remove all trace of them. Either through the use of utility programs or the operating system itself, the original file clusters can be located and any part of the original file that has not yet been overwritten can be regenerated.

Backup tapes, CDs, and DVDs may contain valuable information about the system security structure. For example, in the 1970s and 1980s, system backups on one brand of minicomputer contained the entire directory, complete with every user ID and password *in the clear* on the first tape. Using a simple file copy utility, any user could read these data.

To destroy information on magnetic media, users either must overwrite the medium several times with random data or physically destroy the medium. Degaussers are inadequate unless they meet military specifications, but such units typically cost thousands of dollars.

The problem of readable data is especially troublesome on discarded disk drives or on broken hard disk drives that have been repaired or that are subject to specialized forensic data recovery. Users have received operational, data-laden disk drives as

replacements for their own broken units. Sometimes the replacement disks have not even been reformatted; they contain entire directories of correspondence, customer databases, and proprietary software. Because it is by definition impossible to overwrite data on a defective disk drive, military security specialists routinely destroy defective hard disks using oxyacetylene torches or purpose-built grinders that reduce hard drives into small chunks.

For more information about secure disposal of magnetic and optical data storage media, see Chapter 57 in this *Handbook*.

15.3.4 Spying. Some techniques used by criminal hackers seem to have been lifted directly from spy novels. For example, laser interferometry can reconstitute vibration patterns from reflected infrared laser beams bounced off windows. Users of such equipment can hear and record conversations in rooms that have external windows that vibrate according to the sounds in the room. For more antispy measures related to physical and facilities security, see Chapters 22 and 23 in this *Handbook*.

Hackers surreptitiously steal people's access codes by watching their fingers as they punch in secret sequences. When pay phones were prevalent, *shoulder surfers* would capture telephone calling-card codes, which they sell to organized crime rings. Codes can be stolen by peering over the shoulders of neighboring callers, or by using binoculars, telescopes, and video cameras to track the buttons pressed by their victims.

Shoulder surfing can occur within installations as well. For example, most users of punch-key locks pay no attention to the visibility of their fingers. Whenever punching in a code, users should guard against observation by unauthorized people. In public places, users should stand up close to the keypad. In fixed installations, facilities managers should cover keypads with opaque sleeves allowing unimpeded access but concealing details of the access codes.

Criminals are adept at surfing both wired and wireless network connections, either by cruising around town with a war-driving setup or hanging out in a target-rich environment such as an airport, train station, coffee shop, or hotel lobby. There is more about wireless hacking in Chapter 33 in this *Handbook*.

One particular type of wired connection that should be used with care is the broadband guest-room connection offered by many hotels. Too often, these are set up without proper security measures, enabling a curious or criminally inclined guest to locate machines belonging to other guests (sometimes by merely clicking the Network Neighborhood icon in Windows Explorer). Employees should be instructed not to plug their company laptops into such connections unless they have a properly configured firewall in place and turned on and are using a virtual private network (VPN) to access corporate systems (see Chapter 32 in this *Handbook*). Even some of the most expensive upscale hotels have been found to suffer from this problem, as illustrated in Exhibit 15.1.

15.3.5 Penetration Testing, Toolkits, and Techniques. Verifying and improving the security of systems by attempting to penetrate them is a well-established practice among security professionals and system administrators. However, although some were practicing this technique earlier, it was not openly discussed prior to 1993. That year, Dan Farmer and Wietse Venema released the pioneering paper entitled "Improving the Security of Your Site by Breaking into It."[11] This paper advanced the notion of assessing system security by examining a system through the eyes of a potential intruder. Farmer and Venema showed that scanning for seemingly benign network services can reveal serious weaknesses in any system. Prior to the publication of this important paper, many system administrators were unaware of the extent of vulnerabilities affecting their systems.

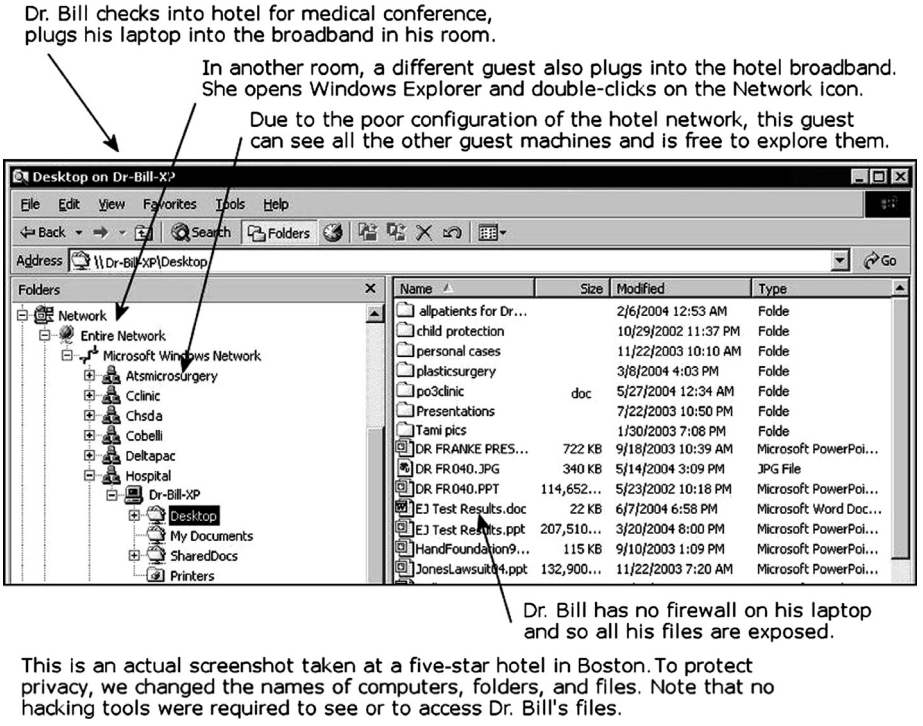


EXHIBIT 15.1 Poorly Configured Hotel Room Internet Connectivity

Farmer and Venema then released a network-testing program called SATAN (Security Analysis Tool for Auditing Networks). Security professionals and system administrators both lauded and were angered by the program. Some system administrators cheered the availability of an all-in-one tool that revealed security holes but did not exploit them. Others questioned the authors' motives in releasing a free and readily available tool that hackers could use to attack networks. While the debate raged on, system administrators and hackers alike began using SATAN to interrogate networks.

15.3.5.1 Common Tools. Since that time, hundreds of penetration toolkits have appeared; they are commonly referred to as *scanners*. Today one can find innumerable freeware tools or invest in one of the commercial tools. Scanners vary in complexity and reliability. However, a majority of the tools employ the same basic functions to test a network: query ports on the target machines and record the response or lack of response.

Used in the proper manner, these tools can be effective in discovering and recording vast amounts of data about a network of computer systems and revealing security holes in the network. Many scanner packages also include packet-sniffing applications, described earlier. Administrators can use this information to reduce the number of systems that can be compromised.

A wide variety of basic network tools may be used in any penetration test. These tools may include mundane programs, such as PING, FINGER, TRACEROUTE, and NSLOOKUP. However, most serious penetration tools make use of an automated vulnerability analysis tool consisting of a series of port queries and a database of known vulnerabilities. Some tools also attempt to exploit identified vulnerabilities in order to eliminate false positives. Once the vulnerabilities have been found, it is

remarkably easy to obtain "exploits" or programs with which to launch an attack against the susceptible machines. All the tools make use of the basic operations of the TCP suite of protocols. Although these protocols will operate on different port numbers, they all share a common structure of a three-way handshake. All TCP protocols look for a connection attempt (connect), a synchronization and acknowledgment exchange (SYN/ACK), various conditions (FLAGS), and a close port request (FIN). Therefore, the operation of port scanners is quite similar. They attempt to find open, or listening, ports on a machine, ask for a connection, and then log the results to a file to be compared to the internal database. The scanner will display the results of the scans by listing the open ports and the services that appear to be running. At this point the various programs differ. Some attempt an in-depth analysis of the possible security holes associated with the ports and services, along with the appropriate security measures to secure them. The more malicious scanners also include automated scripts for exploitation of the vulnerabilities thus revealed.

Of the freeware tools, Nessus, Netcat, and nmap are probably the best known, although there always seems to be a new flavor of the month among both hackers and system administrators. SATAN is still available, as are its spin-offs, SAINT and SARA. A fair amount of skill is required to use these tools as they are fairly sophisticated, and some of the scans can overload a system and cause it to hang or crash.

15.3.5.2 Common Scans. As previously mentioned, most of the scanners/sniffers available will run through the same basic routines in order to develop a picture of a machine as a whole. The purpose is to determine what is running on the machine and what its role is in the network. The next sections describe most basic scans and their results.

15.3.5.2.1 TCP Connect. This is the most basic form of TCP scanning. This system call is provided by the operating system. If the port is listening, the attempt at a connection will proceed. This scan does not require root or supervisor privileges. However, this scan is easily detectable, as many connection and termination requests will be shown in the host's system logs.

15.3.5.2.2 TCP SYN. This scan does not open a full connection and is sometimes referred to as a *half-open scan* because a full handshake never completes. A SYN scan starts by sending a SYN packet. Any open ports should respond with an SYN|ACK. However, the scanner sends an RST (reset) instead of an ACK, which terminates the connection. Fewer systems log this type of scan. Ports that are closed will respond to the initial SYN with an RST instead of an ACK, which reveals that the port is closed.

15.3.5.2.3 Stealth Scans. Also referred to as Stealth FIN, Xmas Tree, or Null scans, the stealth scan is used because some firewalls and intrusion detection systems watch for SYNs to restricted ports. The stealth scan attempts to bypass these systems without creating a log of the attempt. The scan is based on the fact that closed ports should respond to a request with an RST and open ports should just drop the packet without logging the attempt.

15.3.5.2.4 UDP Scans. There are many popular User Datagram Protocol (UDP) holes to exploit, such as an rpcbind hole or a Trojan program, such as cDc Back Orifice, which installs itself on a UDP port. The scanner will send a 0-byte UDP packet to each port. If the host returns a "port unreachable" message, that port is considered closed.

This method can be time consuming because most UNIX hosts limit the rate of Internet Control Message Protocol (ICMP) errors. Some scanners detect the allowable rate on UNIX systems and slow the scan down, so as not to flood the target with messages.

15.3.5.2.5 IP Protocol Scans. This method is used to determine which Internet protocols are supported on a host. Raw IP packets without any protocol header are sent to each specified protocol on the target machine. If an ICMP unreachable message is received, then the protocol is not in use. Otherwise, it assumed to be open. Some hosts (AIX, HP-UX, Digital UNIX) and firewalls may not send protocol unreachable message, so all protocols appear to be open.

15.3.5.2.6 ACK Scan. This advanced method usually is used to map out firewall rule-sets. In particular, it can help determine whether a firewall is stateful or just a simple packet filter that blocks incoming SYN packets. This scan type sends an ACK packet with random-looking acknowledgment/sequence numbers to the ports specified. If an RST comes back, the port is classified as "unfiltered." If nothing comes back, or if an ICMP unreachable is returned, the port is classified as "filtered."

15.3.5.2.7 RPC Scan. This method takes all the TCP/UDP ports found open and then floods them with SunRPC program NULL commands in an attempt to determine whether they are RPC ports and, if so, what program and version number they return.

15.3.5.2.8 FTP Bounce. This scan looks like it is an FTP proxy server within the network (or trusted domain). It could eventually connect to an FTP server behind a firewall. Once the FTP server has been found, scanning of ports normally blocked from the outside can be made from the internal FTP server. Of course, reading and writing to directories can be checked from this server as well.

15.3.5.2.9 Ping Sweeps. This scan uses Ping (ICMP echo request) to find hosts that are up. It can also look for subnet-directed broadcast addresses on the network. These are IP addresses that can be reached externally. Ping sweeps often are used to try to "map" the network as a whole.

15.3.5.2.10 Operating System Fingerprinting. As many security holes are dependent on the operating system of the host, this scan attempts to identify which operating system is running, based on a number of suppositions. It uses various techniques to detect subtleties in the underlying operating system (OS) network stack of the computers being scanned. The data gathered are used to create a "fingerprint" that is compared to the scanner's database of known fingerprints. If an unknown fingerprint is found, attackers can check Websites and newsgroups where information about fingerprints is freely traded, to discover what a particular OS might be. Once the OS has been identified, it is quite easy to find exploits by simply using a search engine on the Web. OS fingerprinting is unnecessary if the OS can be discovered by reading the banners. For example, if one was to telnet to a machine, the response could be:

```
badgny~> telnet abcd.efg.com
Trying 163.143.103.12 ...
Connected to abcd.efg.com
Escape character is '^]'.
HP-UX hpux B.10.01 A 9000/715 (ttyp2)
login:
```

The banner, which was included in the default configuration, simply indicates that the OS is HP-UX. A good system administrator will turn off the banners on all services that have them.

15.3.5.2.11 Reverse Ident Scanning. This scan usually is used to see if a Web server on the network is running as root. If the *identd* daemon is running on the target machine, then a TCP Ident request will cause the daemon to return the username that "owns" the process. Therefore, if this request is sent to port 80 (hex), and the return user is root, then that server can be used for an attack on the system. This scan requires a full TCP connection to the port in question before it will return the username.

15.3.5.2.12 Scanning One's Own Port Configuration. Steve Gibson of Gibson Research Corporation makes available a free port scanner that identifies open ports among the first 1056 TCP ports on any system within less than a minute.[12] Any open port can then be controlled using an appropriate firewall setting. Ideally, all ports on workstations are nonresponsive to pings.

15.3.5.3 Basic Exploits. Using the results of a scanning program, the next logical step for hackers would be to try to exploit the apparent weaknesses in the system. Hackers seek to compromise a machine on the network by getting it to let them run programs or processes at will, at the root level. Once hackers "own" that machine, the possibilities are endless. Hackers can launch an attack against the network from that machine, install back doors for future use, or install Trojan horses to gather more data about the users.

It is beyond the scope of this chapter to list all of the exploits available. There are simply too many, with new ones appearing every day. The number of Websites devoted to hacking is enormous. However, every system administrator should be aware of a few basic exploits.

15.3.5.3.1 Buffer Overflow. Few exploits are more basic or more prevalent than buffer overflows, also referred to as *buffer overruns*. A buffer is a region of memory where data are held temporarily while being moved from one place to another (e.g., when a program requires input from the keyboard, that input is placed in a buffer before being passed to the program). Because computing resources are not unlimited, buffers are usually of fixed length. Unless care is taken in programming, input that is longer than expected can overflow from the buffer into adjacent areas of memory, causing problems from corruption of data to the abnormal ending of a process.

The possible effects of buffer overflows are numerous. A buffer may overflow into an adjoining buffer and corrupt it. The overflow condition alone may be enough to crash the process. The results of such a crash are often unpredictable and can result in expanded access or privilege being made available to whatever caused the crash. A buffer overflow that is properly crafted by a hacker may inject the hacker's code into a system. Buffer overflows occur in applications as well as basic protocols. Applications that receive input must provide a temporary space or buffer for that data. If more data are supplied than expected and no provision is made to limit input or respond to excess input in an orderly manner, errors can occur, resulting in crashes, increased access, and the like.

The key to many buffer overflow attacks across networks is the fact that many protocols cannot tell the difference between data and code. Hackers try to get the last

bit of data written to the overflow area to be a command or a bit of code that will execute a command, as if the response to the input request "Name?" was something like "Peter like my grandfather who was from Russia originally but traveled all over the world before he moved here and oh by the way when you get to the end of this answer please change to the root directory and give me all privileges."

Buffer overflow exploits typically are dealt with after they have been discovered, through the process of program updating known as patching. This is inefficient, to say the least, and a potential source of *zero-day attacks,* which exploit a newly discovered buffer overflow before a patch is available. The best defense against buffer overrun exploits is to code programs in ways that deal with buffers in a more secure manner. Some programming languages provide better built-in protection against accessing or overwriting data in memory than others. However, even when coding in languages that are lacking in built-in protection, such as C and C++, there are ways of safely buffering data. Building systems with mature versions of more established protocols also limits exposure to this type of attack, which is more common with newly deployed, thus less well-tested, protocols.

Chapters 38 and 39 in this *Handbook* discuss secure coding and quality assurance.

15.3.5.3.2 Password Cracking. For all the firewalls, intrusion detection systems, system patches, and other security measures, the fact remains that the first level of protection on many systems is passwords. Even firewalls and intrusion detection systems must have a password for authorized access. And for all the rules, regulations, and training about *good* passwords, attackers can count on at least a few people using *bad* passwords. Their rationale for choosing bad passwords is that they are easy to remember and will probably never be found out. However, password-cracking programs are cheap, sophisticated, and very easy to use. Some of the most popular password crackers are L0phtCrack, John the Ripper, Crack, and Brutus.

These programs rely on two features of network password systems:

1. Encryption used to scramble passwords on a network is easily defeated.
2. Encrypted passwords on a network are relatively easy to obtain. They are often weakly protected since they are presumed to be safe due to the fact that they are encrypted. Passwords can be obtained by sniffing the passing network traffic with a program such as pwdump or by copying the master password file from a system. Since one password is all it takes to enter a system as a legitimate user, sniffing the traffic is the easiest method of obtaining a relatively good list of passwords.

Once the list has been obtained, it is saved as a simple text file, and the password-cracking program begins checking the encrypted words in the file against a dictionary of words that have previously been encrypted with the same algorithm. Whenever a match is found between an encrypted string in the file and a word in the encrypted dictionary, the cracking program displays and records the plaintext of the encrypted dictionary word. Thus the password is revealed.

In addition to checking ordinary dictionary words, some password crackers check for both uppercase and lowercase letters, numbers before and after a word, and numbers used in lieu of vowels within a word. The speed at which these programs operate, even on a basic desktop or laptop computer, is impressive, and it is entirely

possible to obtain cracked passwords within seconds. Indeed, a useful security awareness exercise is to demonstrate such a program to employees: The first passwords to be cracked will be the weakest ones, and that can serve as a warning to users who choose such words.

A good security officer will ensure that passwords on a network are checked regularly with a password cracker or by implementing one of the many strong-password enforcers. Password enforcers augment the password program by comparing the passwords chosen by the user to the rules set by the enforcer.

A simple online tool from Steve Gibson's GRC allows one to calculate the keyspace and estimate cracking time for sample passwords.[13] Users should not enter their actual password but only one using the same rules; for example, if one's actual password were *Dk3*(4n$p2*, one could enter *Eh5&%9g#t8* to arrive at exactly the same analysis. In the example presented here, the calculations result in a keyspace of 6.05×10^{19} with a cracking time of 1 week if a massively parallel array of processors supported 10^{15} guesses per second. The password *evanescent porridge* would take the same processors 14.32 billion centuries to include by brute-force cracking of the 4.50×10^{33} keyspace.

15.3.5.3.3 Rootkits. Rootkits are one of the many tools available to hackers to disguise the fact that a machine has been "rooted." A rootkit is not used to crack into a system but rather to ensure that a cracked system remains available to the intruder. Rootkits are comprised of a suite of utilities that are installed on the victim machine. The utilities start by modifying the most basic and commonly used programs so that suspicious activity is cloaked. For example, a rootkit often changes simple commands such as "ls" (list files). A modified "ls" from a rootkit will not display files or directories that the intruder wants to keep hidden.

Rootkits are extremely difficult to discover since the commands and programs appear to work as before. Often a rootkit is found because something did not "feel right" to the system administrator. Since rootkits vary greatly in the programs they change, one cannot tell which programs have been changed and which have not. Without a cryptographically secure signature of every system binary, an administrator cannot be certain to have found the entire rootkit.

Some of the common utilities included in a rootkit are:

- Trojan horse utilities
- Back doors that allow the hacker to enter the system at will
- Log-wiping utilities that erase the attacker's access record from system log files
- Packet sniffers that capture network traffic for the attacker

15.3.5.3.4 Trojan Code. As described earlier in the context of compromised login procedures, Trojan code is something other than it appears to be. In this case, the Trojans are the changed programs in a rootkit that allow an intruder's tracks to be hidden or allow the program to gather more information as it sits silently in the background. Local programs that are Trojaned often include "chfn," "chsh," "login," and "passwd." In each case, if the rootkit password is entered in the appropriate place, a root shell is spawned.

15.3.5.3.5 Back Doors. Back door utilities often are tied to programs that have been Trojaned. They are used to gain entry to a system when other methods fail. Even if

a system administrator has discovered an intrusion and has changed all the usernames and passwords, there is a good chance that he or she does not know that the back doors exist. To use the back door, the hacker needs only to know the correct port to connect to the compromised machine and to enter a password or command where one is not usually entered.

For example, inetd, the network super daemon, is often Trojaned. The daemon will listen on an unusual port (rfe, port 5002 by default in Rootkit IV for Linux). If the correct password is given after connection, a root shell is spawned and bound to the port.

The function rshd can be similarly Trojaned so that a root shell is spawned when the rootkit password is given as the username and thus rsh [hostname] –l [rootkit password] will obtain access to the compromised machine.

15.3.6 Penetration via Websites. A vast new network territory opened up in the final decade of the twentieth century, partially devoted to commerce and largely driven by attempts to make money from a technology that originally had been developed for military and academic purposes. The relentless growth of this network surpassed 145 million registered domains at the time of writing in May 2013 as reported by the Whois Source (www.whois.sc/internet-statistics), one of the most reliable sources of statistics about the Internet.

Worldwide Internet penetration, measured as the number of Internet users as a percentage of total population, was over one third by June 2012, with North America, surpassing 75 percent.[14] Asia had almost 4 billion users at that time. Not surprisingly, with so many machines in one network and over 2.4 billion users (twice what was reported in the 2009 Fifth Edition of this *Handbook*), this new territory is the primary playground for hackers, from the merely curious to the seriously criminal. The Web presents a target-rich environment for people seeking unauthorized access to other people's information systems. There are several reasons for this; chief among them is the fact that many organizations, both commercial and governmental—including the military—have external, public Websites that are connected, in some way, to internal, private networks. This connection provides a system penetration path that can be exploited in many different ways, as outlined in this section.

For more detailed analysis of Website security, see Chapter 21 in this *Handbook*.

15.3.6.1 Web System Architecture. Standard practice when placing a commercial Website on the Internet is to screen it from hostile activity, typically using a router with access-control lists (ACLs) or a firewall, or both. However, unless the Web site is of the basic, "brochure-ware" kind, which simply exists to provide information on a read-only basis, the site has to allow for user input of data. Input is required for something as simple as a guest book entry or an information request form; more complex applications such as online shopping have more complex input requirements.

A typical method of processing input is the Common Gateway Interface (CGI). This is a standard way for a Web server to pass user input to an application program and to receive a response, which can then be forwarded to the user. For example, when a user fills out a form on a Web page and submits it, the Web server typically passes the form information to a small application. This application processes the data and may send back a confirmation message. This method is named *CGI*, and it forms part of the Hypertext Transfer Protocol (HTTP). Because it is a consistent method, applications written to employ it can be used regardless of the operating system of the server on which it is deployed. Further adding to the popularity of CGI is the fact that it works

with a number of different languages including C, C++, Java, and PERL. The term "CGI" is used generically for Web server code written in any of these languages.

Any system that receives input must provide a path within the system through which that input can flow. This is referred to as an "allowed path." The necessity of allowed paths, combined with the ability to exploit them for penetration attacks, led the system security expert David Brussin (author of Chapter 26 in this *Handbook*) to coin the term "allowed path vulnerabilities" for this class of vulnerability. The two leading categories of exploits that employ allowed paths are input validation exploits, which are akin to buffer overflow exploits and often employed to abuse CGIs, and file system exploits, which abuse the Web server operating system and services running on the server.

Of course, there are other ways to abuse Websites as well. Denial-of-service attacks can be used to inhibit legitimate access and thus compromise availability. (Chapter 18 in this *Handbook* covers DoS attacks.) Not every attack is aimed at further penetration of systems; the Web pages themselves may be the target of attack, as in a defacement, an unauthorized change to Web pages. Defacement often is committed to embarrass the Website owner, publicize a protest message, or enhance the reputation of the criminal hacker who is performing the defacement. But in terms of penetration, the primary goal of attacks on Websites is to compromise internal networks that might be connected to the Web server. Exploitation of allowed path vulnerabilities is probably the most common form of such attacks.

15.3.6.2 Input Validation Exploits. Whenever an allowed path is created to accommodate user input, the possibility exists that it will be abused. Such abuse can lead to unauthorized access to the system. This section describes a range of penetration methods using this approach, all of which somehow employ invalid input, or input that is:

Not expected by the receiving application on the server.

Not "allowed" according to the rules by which the receiving application is operating.

15.3.6.3 Unexpected Input Attacks. How is it possible to submit unexpected input to a server? The answer lies in the architecture of the Internet and the paradoxical nature of the client system that is accessing the server. The typical Web client is a client in name only. It is often a powerful machine, capable of being a server in its own right, and very difficult for any other server to control, due to the inherent peer-to-peer nature of the huge network that is the Internet. All nodes of the Internet are considered hosts. And, of course, many of these hosts are outside the physical control of the organizations hosting those machines that are acting as servers. This fact has serious implications for security.

Unless the server can install tightly controlled application code on the client, and restrict user input to that code, the server must rely on the coding most commonly used to implement Web client-server interaction, Hypertext Markup Language (HTML) and Hypertext Transfer Protocol Daemon (HTTPD). Both are complex and relatively immature. For example, they do not automatically identify the source of the input. Consider an HTML form on a Web page, designed to be presented to a visitor to the Website who fills in the fields and then clicks a button to submit the form. There is nothing on the client side to control the user's input. So instead of entering a first name in the First Name field, a user might enter a long string of random characters. Unless the application processing this data field performs extensive input validation, the effects of

such action can be unpredictable, even more so if the user includes control characters. Similarly, if the Web server itself is not designed to validate page requests, a user might cause problems by submitting a bogus Universal Resource Locator (URL).

The problem is even more severe than this. Unless the Web server application is specifically written to defeat the following abuse, it can be used to cause all sorts of problems that potentially lead to successful penetration. Suppose that, instead of simply filling out a form, the user creates a local copy of the page containing the form, then alters the source code of the form, saves the file, and submits it to the Website instead of the original page. This does not violate the basic protocols of the Web but clearly provides considerable penetration potential. Many Websites are still vulnerable to attacks of this type.

15.3.6.4 Overflow Attacks. As described earlier in the context of hidden form fields, it is possible to gain access to Web servers by supplying more input than expected. Such an attack is possible with any field on a user-submitted form, not just a hidden field. The defense is to build extensive error checking into the application that processes the form.

Overflow attacks, which were described in general terms earlier, also can be directed at applications or services running on the Web server. For example, in June 2001, CERT announced a remotely exploitable buffer overflow in one of the Internet Server Application Programming Interface (ISAPI) extensions installed with most versions of Microsoft Internet Information Server 4.0 and 5.0, specifically the Internet/Indexing Service Application Programming Interface extension, IDQ.DLL. An intruder exploiting this vulnerability may be able to execute arbitrary code in the local system security context, giving the attacker complete control of the victim's system.

15.3.6.5 File-System Exploits. Another category of attack against Websites exploits problems with the file system of the Web server itself. Ever since Web servers started appearing on the Internet, there has been a constant procession of vulnerability announcements arising from file system issues. The main ones are presented here, but the possibility of others appearing is high due to a lack of what David Brussin has called *vulnerability class analysis*. A vulnerability class is a type of problem, such as buffer overflow or file system access control. Developers of Web servers, and many other applications, are often averse to, or resource-constrained from, the elimination of vulnerabilities as a class, being focused instead on the hole-by-hole fixing of specific instances of the vulnerability as they arise. This phenomenon is largely a result of the rapid pace at which the Web has been developed and deployed, driven by powerful commercial forces.

15.3.6.6 Dot Dot, Slash, and Other Characters. Persons responsible for the security of information systems that employ Web servers must be alert for new vulnerabilities. Web server software has proven particularly susceptible to certain categories of vulnerability that tend to recur in new versions. Whenever these vulnerabilities are discovered, attackers quickly exploit them. Typically, software vendors issue patches to solve the problem, but systems remain susceptible until patched, and attackers use automated tools to scan the Internet for servers that are still susceptible. For example, in April 2001, a flaw was discovered in versions of Microsoft Internet Information Server (IIS) in use at that time. This flaw made it possible for remote users to list directory contents, view files, delete files, and execute arbitrary commands.

In other words, if a Web server running IIS was connected to the Internet, anybody using the Internet could potentially copy, move, or delete files on the Web server. With this level of access, it was possible to use the Web server to gain access to connected networks unless strong internetwork access controls were in place. In alerts that were issued to warn users of this software, exploitation of this vulnerability was described as "trivial." In fact, a large number of sites were penetrated because of it, and many suffered defacement (i.e., unauthorized changes to the appearance of their Websites). In other cases, attackers downloaded sensitive customer data and uploaded and installed back door software.

Particularly worrying about this vulnerability was the fact that it was essentially a recurrence of the so-called dot dot directory traversal attack, which was possible on a lot of early Web servers. These servers would, upon request, read ".." directories in URLs, which are unpublished parent directories of the published directory. Thus, attackers were able to back out to the Web root directory and then to other parts of the server's directory structure. This technique basically allowed attackers to navigate the file system at will. Many Web servers, including IIS, began to incorporate security measures to prevent the "dot dot" attack, denying all queries to URLs that contain too many leading slashes or ".." characters.

The vulnerability published in April 2001 involved bypassing these restrictions by simply substituting a Unicode translation of a "/" or " ∴" Attackers found that, by appending the ".." and a Unicode slash or backslash after a virtual directory with execute permissions, it was possible to execute arbitrary commands. Attackers could execute any command via a specially crafted HTTP query. The frequency with which "old" vulnerabilities reappear in new software should serve as a warning to information security professionals not to assume that "new" is the same as "improved." Indeed, all software needs to be treated with a healthy degree of skepticism and a fair amount of heavy testing prior to deployment.

15.3.6.7 Metacharacters. Dots and slashes used in field system references are closely related to metacharacters, which also can be used to attack Web systems. A metacharacter is a special character in a program or data field that provides information about other characters, for example, how to process the characters that follow the metacharacter. Users of DOS or UNIX are probably familiar with the wildcard character, a metacharacter that can represent either any one character or any string of characters. If used inappropriately, for example, in user-supplied data, metacharacters can cause errors that result in unintended consequences, including privileged access.

15.3.6.8 Server-Side Includes. Server-side includes (SSIs) are special commands in HTML that the Web server executes as it parses an HTML file. SSIs were developed originally to make it easier to include a common file, called an *include file*, inside many different files; examples include files containing a logo or text files consisting of the page, date, author, and so on. This capability was expanded to enable server information, such as the date and time, to be included automatically in a file. Eventually, several different types of include commands were provided on Web servers: *config, include, echo, fsize, flastmod, exec*. The last of these, *exec*, is quite powerful, but it is also a security risk, as it gives to the user of the Web client permission to execute code. A number of attacks are possible when *exec* is permitted within an inadequately protected directory.

Analogous to SSIs are ASPs (Active Server Pages) and JavaServer Pages, as well as Hypertext Processors (PHP). All are technologies that facilitate dynamic page building

and allow execution of codelike instructions within an HTML page. All should be employed with special attention to security implications and strict adherence to secure Web application programming methods. Security professionals will note that many of the weaknesses in these newer technologies are simply old exploits reborn. Anything that offers "greater interactive functionality" may increase the likelihood of new vulnerabilities simply because, in the absence of rigorous software quality assurance, changes are so often associated with new bugs.

15.3.7 Role of Malware and Botnets. Several new twists on system penetration via Websites and Internet email have emerged in recent years, starting with the evolution of computer viruses and worms. Some worms and viruses are designed to install Trojan code. The term "malware," derived from "malicious" and "software," is now widely used to refer to the entire category of software coded with malicious intent, including viruses, worms, and Trojans. In fact, viruses and worms, which are dealt with in more detail in Chapter 16 in this *Handbook*, are themselves a form of system penetration; after all, the creators of virus and worm code are getting their code onto systems that they are not authorized to use, so one may consider those machines to have been penetrated. Some criminal hackers have combined different elements of malware to penetrate computers used to surf the Web. Those compromised machines are then used, in turn, to spread malware and compromise additional machines. The goal may be gathering of user names and passwords (helpful for yet more system penetrations) or financial data used to perpetrate fraud and identity theft.

The two main components of this penetration strategy are drive-by downloads and botnets. A drive-by download attempts to compromise machines used to visit a malicious Website, taking advantage of either user gullibility or vulnerabilities in their Web browsers to install, without explicit permission, unsolicited code, typically a Trojan of some sort. A botnet is a collection of bots, host computers that can be controlled remotely (i.e., ro*bot*ically) through Trojan code installed on those machines, either via a drive-by download or other means, such as a virus or worm. Here is how researchers at Google described the phenomenon in a landmark 2007 report:

> [C]omputer users have become the target of an underground economy that infects hosts with malware or adware for financial gain. Unfortunately, even a single visit to an infected Website enables the attacker to detect vulnerabilities in the user's applications and force the download of a multitude of malware binaries. Frequently, this malware allows the adversary to gain full control of the compromised systems leading to the ex-filtration of sensitive information or installation of utilities that facilitate remote control of the host.[15]

What makes the Google report a landmark is not the existence of drive-by exploits, which have been on the rise for several years, but their prevalence. Because Google maintains a massive repository of Web pages in order to operate its search engine, it is in a fairly unique position when it comes to analyzing the content of the Web as a whole. The solidly researched finding that at least 1 in 10 of all Web pages contained some form of malware should be a wakeup call to IT departments everywhere. Are your users surfing to Facebook pages? Are they aware that something as seemingly harmless as visiting their favorite singer's page on the Web might cause malicious code to be downloaded onto their computer from a server in China?

That sort of attack started to become commonplace in 2007, as documented by Roger Thompson of Exploit Prevention Labs, with Alicia Keys being one of a number of artists targeted by criminal hackers. One exploit used in these attacks installed a

proxy network bot, known as a *flux bot*. The goal of a flux bot is to obscure the location of phishing sites by using constantly changing proxy servers, thus making it harder for banks and other institutions targeted by the phishing scam to shut it down. (There is more about phishing scams in Chapter 20 in this *Handbook.*)

According to the Google report, there are four main methods by which Web pages are turned into malware infection vectors: advertising, third-party widgets, user-contributed content, and Web server security. This implies that company Websites not only need to be firmly secured, but that all advertising and user-supplied content should be validated, as should any widgets that are employed or distributed by the site.

15.3.8 Sophisticated Attackers. In recent years several groups of attackers have been very successful using a very different approach to compromise security systems than has been used historically. These groups have been termed "Advanced Persistent Threats" by the United States Air Force in 2006. While many of these groups are nation state–sponsored these same techniques are being employed by groups with purely financial gain as well.

It is important to understand up front that these attacks, regardless of specific source group, are targeted. In the case of the nation state–sponsored groups, the goal is information theft. This information theft often takes the form of email, but also can range much deeper to research and development information or other intellectual property. The financially motivated groups focus on financial institutions for the purposes of financial gain.

As of this writing, there are many known groups active in the world. Each of these groups uses a little different specific components in their overall process, but the overall process is the same and can be broken down into some distinct stages.

The first stage is reconnaissance, planning, and preparation. This is where they select specific individuals within the target organization. They also plan out their attack based on the results of the reconnaissance. Finally they set up specific command and control servers to be used against the target and build the malicious software to be used as part of the attack.

The second stage is obtaining foothold and persistence. In this stage their goal is to get onto the target network. This is the point they deploy their attack (most often through spear-phishes) and gain remote access to the target network. Since gaining a foothold is one of the more challenging parts, they also seek to ensure they can maintain that foothold on the network.

The third stage is mapping the internal network and locating the desired data. At this point, they seek to understand the target environment and find their ultimate target of either specific data or systems on which they can perpetrate their financial impact.

The fourth and final stage is exfiltration or exploitation. This is the point at which they transmit the collected information outside of the target environment (typically nation-state actor goal) or execute their financial attack (in the case of the cybercrime actors).

15.3.8.1 Stage One—Reconnaissance, Planning, and Preparation. The attackers often begin with selecting a handful of individuals at the target organization. They employ open source intelligence sources like the company Website, Facebook, and LinkedIn to gather information and identify recipients at the target organization. Other sources, such as press releases from companies and news articles about those companies, provide a rich resource for attackers to zero in on individuals close to the ultimate target of the attackers.

Once individuals have been identified an appropriate delivery can be created. The most common delivery means is through spear-phishes. A spear-phish is a targeted phish (as compared to a normal phish that is sent to thousands or millions of recipients). These spear-phishes are often simple and straightforward but can be quite sophisticated.

While spear-phishing is the most common means of delivery attackers will also compromise Websites commonly used by their targets as a means of delivery. This method of delivery is referred to as watering hole attacks. The term *watering hole* refers to the hunting method whereby a hunter waits for their prey at a watering hole, knowing that eventually the target will come for a drink. Watering hole attacks work by modifying a compromised Website's code so that when users view the Website an exploit is delivered to the victim browser. This exploit is in turn used to install malicious code on the system.

Least commonly, the sophisticated attackers have been found to use malicious code placed on USB keys. These are either left where target users will find them or even mailed directly to the targets. Exploit code is used on the USB device so that when users access the device the exploit code will install malicious code onto the user system.

Also as part of stage one the attackers will configure C2 (command and control) servers to be used over the course of the attack. These C2 servers are often compromised hosts at legitimate organizations but can also be hosting servers obtained legitimately. In particular, attackers prefer servers at places such as large universities, as these are unlikely to be blocked by common prevention mechanisms like Web proxy filters commonly employed by organizations today.

The final major aspect of the preparation is the creation of the backdoors to be employed against the target organization. The attackers generally have several variants of their backdoor with slightly different characteristics. Some groups will even use commonly available RATs (Remote Access Trojans) such as Poison Ivy and configure them to use the appropriate C2s. Many of the attackers have tools that will weaponize legitimate PDF or office documents. Weaponization is the process of turning a legitimate file into a malicious one. If, for instance, the individuals share a particular industry in common, the attackers might find a relevant industry conference in the near future and download an actual agenda PDF file from the conference Website. They might then weaponize that PDF or simply use information from the file to add to their spear-phish to add realism and improve their chances for the targeted users to fall for the spear-phish.

15.3.8.2 Stage Two—Foothold and Persistence. The end result of stage one is an RAT being installed on one or more target systems. This Trojan provides covert backdoor access to the victim systems. Now that the attacker has a foothold on the target network, they want to maintain that access. Attackers want access from more than a single system, since a single host might not be running when they want access. The attackers also understand that eventually the RAT will be found or removed. The attackers also want to increase their access to ensure they can get to whatever their final objective is.

To accomplish all of this, the attackers will usually follow a consistent process. They start with dumping the password hashes from the local host cache. By default, Microsoft Windows (the predominant operating system in use by users at most organizations) keeps a cache of the last 10 user IDs and passwords used to log into a computer. This is done so that if the computer is not connected to the corporate environment and thus can't access the domain controllers, the user will still have the ability to log into the system locally. Tools like pwdump, Windows credential editor, and fgdump

will display the user IDs and passwords in the local system cache. These can then be cracked using rainbow tables or cracking software back at the attacker's location.

The second step is to look at the host network connections to determine what computers the host is communicating with. The logic behind this is simple—any credentials in the local computer will very likely work on other hosts the system is interacting with. Remember that the purpose of user IDs is to provide authentication to other systems for the purpose of using them. Using the cracked password for the locally cached user IDs, the attacker connects to the subsequent systems and proceeds to extract the user IDs and password hashes from that host. The process of moving from host to host like this is termed *lateral movement*. It is especially important to note that the attackers are not using malware to accomplish this access. They are using legitimate user IDs and passwords. Ultimately, the attacker will continue this process until they find a system with an administrative account that will give them access across the entire environment. It is rare for them to have to traverse more than six to ten systems in order to find one with a network or domain administrator account cached. Once they have an administrative account they are free to use it to move throughout the network.

In addition to obtaining local user IDs and passwords from system caches, the attackers also map the target environment at this stage. Mapping the network is done through simple querying of Active Directory. The simplest way for the attackers to gain information about the environment is with a series of simple commands:

```
net group "domain computers" /domain
net group "domain users" /domain
net group "domain controllers" /domain
net group "domain admins" /domain
```

If you aren't familiar with those commands, you should try them. The only requirement is to do them from a computer that is a member of the local domain. No administrative credentials are necessary. The result is a list of all computers, users, domain controllers, and domain admin accounts in the environment. This information is returned immediately and is far more useful for attackers' purposes than running a tool like nmap. A further benefit is that the use of these commands is very, very difficult to detect as compared to nmap or similar tools that are very noisy. When the attackers want to gain more thorough information about an environment they will typically use dsquery and dsget. These are two command line tools provided by Microsoft to query active directory. Here are the contents of a batch script that was recovered from a recent incursion showing the use of these tools.

```
dsquery user -limit 0 | dsget user -samid -display -title
   -email -dept -office -company -c >1
dsquery user -limit 0 | dsget user -samid -pwdneverexpires
   -acctexpires -loscr -profile -hmdir -hmdrv -c >2
dsquery user -limit 0 | dsget user -c > 3
dsquery group -limit 0 |dsget group -c >g
dsquery subnet -limit 0 |dsget subnet -c >sn
dsquery site -limit 0 |dsget site -c >s
dsquery computer -limit 0 |dsget computer -c>c
dsquery ou -limit 0 | dsget ou -c>o
```

Attackers understand that users leave organizations and change roles. To ensure they have plenty of user IDs at their disposal it is a high priority for them to dump the entire domain credentials. This usually occurs within minutes of them achieving an administrative account with sufficient credentials. Their tool of choice for this is the

well-known pwdump. They simply run pwdump against one of the domain controllers using a recovered domain administrator account. Once the hashes are obtained, they can crack them at their leisure back at their location.

Another priority of the attackers is to install additional mechanisms for gaining access to the environment. To accomplish this they will use a variety of methods. If a company employs a single-factor VPN, then the attackers have unlimited access to the environment using the stolen user credentials. They will also install several additional backdoors. These are distributed throughout the environment to make it difficult for security personnel to find them all. The attackers understand that malware is the easiest thing to spot in an environment, however, and so they will only install copies on a handful of the overall hosts they compromise. As a further precaution to detection, they will usually configure the extra backdoors to sleep for extended periods of time. This tactic results in it being very challenging for security staff to find and eliminate all copies of their backdoors. Finally, in the last couple years, they have been observed also deploying Web shells in corporate DMZs as a further level of backup access to an environment. Web shells were very popular with attackers in the late 1990s and early 2000s but have been rare for the last several years, and thus are not commonly looked for. The Web shells being employed by the attackers are very small and do nothing until remotely accessed, making them very difficult to uncover as well.

As if the use of so many techniques in order to maintain access to a company was not enough, most groups will also monitor their installed backdoors in real time. If they observe an organization actively removing them they will quickly jump onto compromised hosts and install entirely different backdoor versions with very different characteristics so as to maintain access and remain undetected. It is important to realize that the majority of the backdoors are proprietary, not commodity, and very rarely found with standard antivirus. The net result for the attackers is a long-time presence on an environment, during which they can move around and access whatever they desire.

15.3.8.3 Stage Three—Information Theft. The final stage and aspect of the advanced attackers is where they find and exfiltrate the information they see or execute the financial transfers in the case of the cybercrime-motivated actors. All of the efforts of the attackers are for the ultimate purpose of stealing something.

After establishing long-term persistence in an environment, the attackers begin moving from host to host looking for the information they seek. Mostly this is done by command line, but some groups will also tunnel RDP protocol through their C2 infrastructure. If an organization has some other form of remote control, such as VNC, the attackers will gladly utilize that as well with the legitimate credentials they compromised. Through the simple expedient of the Windows 'dir' command, they look in the documents of each host, seeking the data they desire. This is facilitated by the deep organizational knowledge they acquired in earlier stages.

Often the majority of their theft is simply e-mail. When you consider the amount of business activity details that are conducted with e-mail, you realize an attacker can gain incredible competitive advantage from information contained in e-mail.

When the attackers find the data they seek, they will compress and encrypt it, then transmit it back to the C2 infrastructure. This process is known as *exfiltration*. Since sophisticated attackers understand organizations have logging measures in place, the transfer is done to intermediary systems, rather than directly to their location. Subsequently they can move it from the intermediary systems to the final destination without risk of the compromised organization being able to trace it.

Most often RAR is used to package the data for exfiltration. RAR is particularly useful for attackers because of its unique characteristic among compression formats: Data in an RAR file can be recovered even if the RAR is incomplete. If exfiltration is interrupted midtransmission, the attackers will still be able to extract whatever portion of the data made it out. Encryption is used to make it more difficult for a compromised company to figure out what data was stolen. Most groups also prefer to use a media file extension for their RAR files as a further means of obfuscation. The renamed files may have a .mpg or .avi extension but they are simply normal encrypted RAR files.

For transmission of the data the attackers typically use Secure Sockets Layer (SSL). This tool, combined with the compressed and encrypted contents, makes it very difficult for companies to detect that anything malicious is occurring. The sophisticated attackers are so successful overall because the majority of their activities are done by hiding in the large volume and noise of legitimate activity. Consequently, they've all been observed using services like dropbox, box.net, and other cloud providers for exfiltration as well, given the significant adoption rate of these types of services legitimately by our users. Detecting malicious activity to a cloud provider from legitimate activity is very difficult indeed.

15.4 POLITICAL AND LEGAL ISSUES.

Thanks to the World Wide Web, the Internet has become a self-documenting phenomenon. One can use the Internet to find out everything one wants to know about the Internet, including how to penetrate information systems that employ Internet technology. However, the penetration information available on the Internet is not restricted to Internet systems, and the Internet is not the only source of penetration information. Furthermore, the very availability of penetration information is fraught with political and legal issues. These are discussed briefly in this final section of the chapter.

15.4.1 Exchange of System Penetration Information.

The sharing of system penetration information—that is, information that could facilitate illegal penetration of an information system—is the subject of a long, heated, and ongoing debate. This debate encompasses both practical and ethical aspects of the issue. Although complete coverage is not possible within the confines of this chapter, we do review the question of full disclosure, along with some of the sources for penetration information.

15.4.2 Full Disclosure.

How should we handle known vulnerabilities and potentially damaging computer viruses? Should we publish full details, conceal some details, or suppress any publication that would allow exploitation until patches or updates are available from manufacturers? Is there a case for handling some viruses and exploits differently from others?

Over the last two decades, formal venues for full disclosure of system or network vulnerabilities and exploits have evolved (e.g., BugTraq). Support for full disclosure of such details, down to the source code or script level, from professional, honest security experts (the Good Guys) is based on subsets of several key beliefs:

- The Bad Guys know about the vulnerabilities anyway.
- If they do not know about it already, they will soon with or without the posted details.
- Knowing the details helps the Good Guys more than the Bad Guys.

- Effective security cannot be based on obscurity.
- Making vulnerabilities public may force vendors to improve the security of their products.

Since those who would use vulnerabilities and exploits for gain or harm often learn of them before system administrators and security experts, it makes sense—so the argument goes—for the Good Guys to spread the knowledge where it can do some good. One might also argue that, because cryptographic techniques are routinely exposed to public scrutiny by experts to detect vulnerabilities and to avoid future failures, other aspects of security should also be made public.

As for putting pressure on manufacturers, one colleague describes an incident that illustrates the frustrations that sometimes underlie full disclosure. He informed an important software supplier of a major security vulnerability. The product manager ignored him for a month. At that point, patience gone, the colleague informed the product manager that he had exactly one more day in which to produce a patch; otherwise, he said, he would publish the hole in full in the appropriate Usenet group. A patch was forthcoming within one hour.

Why would anyone object to full disclosure of detailed viral code and exploits? The arguments are that:

- Nobody except researchers needs to know the details of viruses or even of specific exploits.
- Publishing in full gives credibility to ill-intentioned Bad Guys who do the same.
- Full disclosure makes impressionable youths more susceptible to the view that illegal computer abuse is acceptable.

How, exactly, does publishing the details of a new virus help system administrators? In one view, such details should be exchanged only among colleagues who have developed trust in each other's integrity and who have the technical competence to provide fixes. For example, a "zoo" of computer viruses serves this function, with access limited to legitimate virus researchers who sign a code of ethics that forbids casually distributing viruses to anyone who wants samples. Opponents of this stance see the attitude as arrogant and elitist. Furthermore, some virus and worm code is written in a form that is easily read by anyone who receives a copy (a large population, considering that in 1999, the Melissa virus is thought to have infected over a million computers within a matter of days).

There is a danger in publishing exploits where any person with access to the Web can use them for automated attacks on Websites. This gives naïve people the impression that it is okay to publish any attack code, regardless of consequences. (Note that the consequences are often relatively minor—the author of the Melissa virus, which is estimated to have caused at least $80 million in damages, served only 20 months in a federal prison and paid a fine of only $5,000.) What is the difference, then, between publishing vulnerabilities or exploits and actually creating attack tools? Was creating and publishing BackOrifice a morally neutral or even a useful act? BackOrifice is a tool that is explicitly designed to install itself surreptitiously on systems and then hide in memory, using stealth techniques modeled on what some viruses use. Is this a contribution to security?

There are no simple or uncontested answers to these questions, but in some cases technology has answered them for us. Before 1995, when macro viruses first appeared

"in the wild," most viruses were written in assembly language or machine code. This limited the number of people who could read them or understand them to persons familiar with assembly language and machine code. However, anyone who receives a macro virus has the text editor tools needed to read its code. The tools required to develop and test macro viruses are provided in mainstream applications such as Microsoft Word, which has tens of millions of users worldwide, a significant percentage of whom have, or can easily acquire, the rudimentary programming skills required to develop their own viruses. The rapid spread of the original Word Concept virus in the summer of 1995 pretty much ensured that anyone who wanted a copy of it could get one (and many of those who did not want a copy got one anyway). The idea of keeping the content of the virus secret was a nonstarter.

Faced with vendor inability or unwillingness to fix security vulnerabilities in a timely manner, some users and experts can be expected to turn to full disclosure of security information, even though many of them may deplore using such tactics. However, they will always run the risk of making the wrong call when it comes to the effect of such disclosures, which are inherently unpredictable. Indeed, the release of the Word Concept virus may have been motivated by a desire to make Microsoft change the way its office applications handle macros. Other office applications, such as Word Perfect and Lotus 1-2-3, were designed to keep macro code separate from document content, making a malicious document much harder to create.

15.4.3 Sources. There are many sources for information about how to penetrate systems. The motives behind these sources range from highly ethical to downright criminal. Chapter 74 in this *Handbook* discusses training and certification in penetration testing for legitimate, authorized purposes.

15.4.3.1 Online Sources. There is no small irony in the fact that much of what a person needs to know about how to penetrate information systems is made available by information systems. Fortunately, the inverse is also true, as one of this chapter's authors has observed: "The best weapon with which to protect information is information."[16] For this reason, security professionals need to know what sources of penetration information are available.

Today there are thousands of Websites that

- Document security holes in different versions of operating systems
- Distribute hacking tools and discuss how to use them
- Catalog default credentials for network hardware
- Teach malicious code writing, including how to make viruses and worms
- List license codes to enable activation of pirated software
- Buy, sell, and trade system access codes, stolen credit cards, stolen identity data, and networks of compromised hosts (botnets)
- Provide a forum and meeting place for those seeking to penetrate systems

These sites have assumed the mantle of earlier online communication channels, such as bulletin boards and Usenet groups, where legitimate sharing of information occurred alongside the exchange of illegal information, pirated code, and so on. Some Websites are moderated and so maintain certain ethical standards. Others follow the Internet tradition of "anything goes." System administrators and employees should

never participate in discussions on these Websites with company email addresses. In a popular strategy, hackers wait for platform-specific vulnerabilities to be announced, then search the Internet for messages from people using that platform, look at their email addresses to see where they work, and attack systems at those companies in the hope that patches are not yet installed.

15.4.3.2 Publications. Over the years many publications have specialized in hacking. Often these provided information on how to penetrate systems. Some publications were primarily electronic, such as *Phrack,* while others have been print based, such as *2600*, which is now widely distributed through conventional magazine channels as well as through paid subscriptions.

15.4.3.3 Hacker Support Groups. Numerous groups of people exist to share hacking information, ranging from relatively stable entities with their own publications (e.g., *2600* and *cDc*) to annual conventions with thousands of participants (e.g., DefCon). Although some members of these groups may have committed criminal acts, and some participants at hacker conventions actually have been convicted of such and served time, there is usually a diverse mix of elements with different motivations in these groups and meetings. DefCon, for example, draws not only people who openly advocate unauthorized security testing of other people's systems and networks, but also law enforcement personnel and legitimate security experts. Some participants go to DefCon specifically to convince young people not to break laws while trying to learn about security.

Many security professionals would prefer that the line between white-hat hacking and black-hat hacking be clearer and more sharply enforced; however, some companies overlook past transgressions in order to gain the perceived value of the hackers' technical security expertise. Indeed, in recent years, investors have even cooperated with groups of hackers in founding security consulting companies, complete with some employees who continue to use hacker handles. The fact remains that some of the best technical training in the field is provided by people who gained their expertise in criminal (or quasi-criminal) hacking, but who now help defend against such activity.

15.4.4 Future of Penetration. Trends over the last 15 years strongly suggest that attempts to penetrate information systems will not decrease any time soon. These factors have been in play for some time:

- The declining cost of, and increased access to, penetration technology—from software and hardware used to crack passwords and encryption to eavesdropping and interception devices
- The continuing practice of fielding inadequately tested systems, built with immature technology and with insufficient attention to security
- The increased availability of automated hacking tools with easy-to-use interfaces
- The continuing allure, and portrayal in popular culture, of hacking as a "cool" activity, without adequate reflection on its legality or consideration of its morality

Additionally, these factors have emerged strongly in recent years:

- The very real opportunity to make money from penetrating systems, given a thriving market in purloined personal data, compromised hosts (bots), and exploits

- The increased interest and involvement of organized crime in system penetration
- The rise of transnational terrorist organizations that are increasingly computer-literate and may be inclined to penetrate systems belonging to entities or countries to which they are opposed[17]

Although some companies and government agencies are actively pursuing improved responses to these threats, others are not. Through lack of concern, resources, or time to address these trends, many entities are at increasing risk. Each new technology brings new threats, but new threats typically are discounted as scaremongering by vendors who offer defenses against them. For many companies and government agencies, the enthusiasm to reap the benefits of new technology overrules the warnings about risks inherent in its deployment. When those risks finally manifest themselves in ways that threaten the organization, the reaction is usually to buy a technical fix, while the root causes of vulnerability, namely human behavior and employee awareness of security issues, fail to receive the attention and resources they deserve.

New applications of computer technology, such as implanted medical devices and control systems for automobiles and autonomous vehicles, are providing fertile ground for experimentation by research scientists and criminal hackers, who are finding many vulnerabilities.[18]

15.5 SUMMARY

- Penetration of information systems is possible by means of a wide range of methods, some of which are very hard to defend against.
- Those responsible for securing systems have to defend against this wide range of penetration methods.
- Making sure all defenses against all attacks are effective all of the times is a lot harder than finding a single point of failure within those defenses.
- Although all systems do not need to defend against all types of attack equally, the cost of even the more exotic attack strategies is constantly falling, expanding the range of possible attackers.
- The cheapest and most effective attacks are often nontechnical, exploiting human frailty rather than weaknesses in the technology.
- Experienced criminal hackers tend to favor the nontechnical attack over the technical; and the best defense, employee awareness, is also nontechnical.
- Systems can be attacked at the client, at the server, or at the connection between the two.
- Both wired and wireless systems are highly susceptible to eavesdropping and interception.
- Many systems today are built with immature and insecure technology, making them susceptible to a wide range of attacks.
- New attacks come to light with alarming but predictable regularity.
- Many of these new attacks are old attacks reborn, due to a lack of vulnerability class analysis. As a result of economic pressures, faulty reasoning, and insufficient desire for security, vulnerabilities are fixed one instance at a time rather than one class at a time.

- Allowed path attacks against Websites are consistently the most effective strategy for system penetration whenever a system is Web connected or Web enabled.
- Penetration testing, both by internal staff and by objective external experts, always should precede system deployment.
- Given the inevitability of penetration attempts and the high probability of their eventual success, systems should be designed to survive attacks, limiting the scope of compromise from any single point of failure.
- Penetration of systems will continue to fascinate the curious and tempt them to break the law by illegally accessing systems. The potential gains from system penetration, in terms of money, power, competitive advantage, and notoriety, will continue to motivate those to whom laws and morality are not effective deterrents.
- Penetration of systems will become increasingly automated and simplified, further widening the range of possible attackers.
- Human nature, not technology, is the key to defense against penetration attempts. Only by raising society's ethical standards and educating employees to understand the willingness of others to behave unethically can the occurrence of criminal hacking into information systems be significantly reduced.

15.6 FURTHER READING

Websites

CERIAS Hotlist: www.cerias.purdue.edu//hotlist

INFOSEC and INFOWAR Portal: www.infowar.com

SANS InfoSec Reading Room—Penetration Testing: www.sans.org/reading_room/whitepapers/testing

SearchSecurity.com: http://searchsecurity.techtarget.com

SecurityFocus: www.securityfocus.com

Web Application Security Consortium: www.webappsec.org

BooksAllen, Lee. *Advanced Penetration Testing for Highly-Secured Environments: The Ultimate Security Guide*. Packt Publishing, 2012.

Chappell, Laura. *Wireshark® 101: Essential Skills for Network Analysis*. Laura Chappell University, 2013.

Engebretson, Patrick. *The Basics of Hacking and Penetration Testing: Ethical Hacking and Penetration Testing Made Easy*. Syngress, 2011.

Faircloth, Jeremy. *Penetration Tester's Open Source Toolkit*, 3rd ed. Syngress, 2011.

Fialka, J. J. *War by Other Means: Economic Espionage in America*. New York: W. W. Norton, 1999.

Goodell, J. *The Cyberthief and the Samurai: The True Story of Kevin Mitnick—and the Man Who Hunted Him Down*. New York: Dell, 1996.

Kennedy, David. Jim O'Gorman, Devon Kearns, and Mati Aharoni. *Metasploit: The Penetration Tester's Guide*. No Starch Press, 2011.

Litchfield, D., Anley, C., Heasman, J., and Grindlay, B. *The Database Hacker's Handbook: Defending Database Servers*. Hoboken, NJ: John Wiley & Sons, 2005.

McClure, S., Scambray, J., and Kurtz, G. *Hacking Exposed*: Network Security Secrets & Solutions, 7th ed. New York: McGraw-Hill Osborne Media, 2012.

McGraw, G. *Software Security: Building Security In*. New York: Addison-Wesley Professional, 2006.

O'Connor, T. J. *Violent Python: A Cookbook for Hackers, Forensic Analysts, Penetration Testers and Security Engineers*. Syngress, 2012.

Scambray, J., Shema, M., and Sima, C. *Hacking Exposed: Web Applications*, 2nd ed. New York: McGraw-Hill Osborne, 2006.

Schwartau, W. *Pearl Harbor Dot Com*. Seminole, FL: InterPact Press, 2001.

Shimomura, T., and J. Markoff. *Takedown: The Pursuit and Capture of Kevin Mitnick, America's Most Wanted Computer Outlaw—by the Man Who Did It*. New York: Hyperion, 1996.

Slatalla, M., and J. Quittner. *Masters of Deception: The Gang that Ruled Cyberspace*. New York: HarperCollins, 1995.

Sterling, B. *The Hacker Crackdown: Law and Disorder on the Electronic Frontier*. New York: Bantam Doubleday Dell, 1992.

Stoll, C. *The Cuckoo's Egg: Tracking a Spy through the Maze of Computer Espionage*. New York: Pocket Books/Simon & Schuster, 1989.

Stuttard, D., and Pinto, M. *The Web Application Hacker's Handbook: Discovering and Exploiting Security Flaws*. Hoboken, NJ: John Wiley & Sons, 2007.

15.7 NOTES

1. "Report on the Existence of a Global System for Intercepting Private and Commercial Communications (ECHELON Interception System)," PE 305.391, July 11, 2001, http://cryptome.org/echelon-ep-fin.htm
2. S. Cobb, *The Stephen Cobb Guide to PC & LAN Security* (New York: McGraw-Hill, 1992).
3. Robert Lemos, "Mitnick Teaches 'Social Engineering'," News.com, published on ZDNet News, July 17, 2000, http://news.zdnet.com/2100-9595_22-522261.html (URL inactive)
4. "Defector Smuggled Out Copies of the 'Crown Jewels' of Soviet Espionage," *The Times (London)*, September 12, 1999.
5. Bryan Betts, "Optical Nets Easier to Hack than Copper," *PCWorld*. April 27, 2007, www.pcworld.com/article/131306/article.html
6. Madeleine Acey, "'Scrambler Software' Will Protect Phone Calls from Prying Ears," *CNN | Edge of Discovery*, August 17, 2012, www.cnn.com/2012/08/11/tech/silent-circle-encryption
7. AT&T Press Release, Dallas, TX, December 6, 2007, www.att.com/rss
8. Wim van Eck, "Electromagnetic Radiation from Video Display Units: An Eavedropping Risk?" Cryptome, 1985, http://cryptome.org/emr.pdf
9. NSA CSS, "TEMPEST Level I," NSA, January 15, 2009, www.nsa.gov/applications/ia/tempest/TEMPESTLevel1.cfm
10. Distributed Rainbow Table Project, "Free Rainbow Tables," Free Rainbow Tables, May 8, 2013, https://www.freerainbowtables.com
11. Dan Farmer and Wietse Venema, "Improving the Security of Your Site by Breaking into It," available on http://nsi.org/library/compsec/farmer.txt
12. Steve Gibson, *Shields UP!!* 2013, https://www.grc.com/x/ne.dll?bh0bkyd2
13. S. Gibson, "How Big Is Your Haystack … and How Well Hidden Is YOUR Needle?" March 28, 2012, https://www.grc.com/haystack.htm
14. "According to Internet World Stats," www.internetworldstats.com/stats.htm

15. "The Ghost in the Browser: Analysis of Web-based Malware," www.usenix.org/events/hotbots07/tech/full_papers/provos/provos.pdf
16. Cobb, *Stephen Cobb Guide*.
17. C. Wilson, "Botnets, Cybercrime, and Cyberterrorism: Vulnerabilities and Policy Issues for Congress," Congressional Research Service Report for Congress, RL32114, 2007, at http://tinyurl.com/2y995r
18. A. Rubin, "*All Your Devices Can Be Hacked*," TED, October 2011, www.ted.com/talks/avi_rubin_all_your_devices_can_be_hacked.html

CHAPTER 16

MALICIOUS CODE

Robert Guess and Eric Salveggio

16.1 INTRODUCTION **16·1**

16.2 MALICIOUS CODE THREAT MODEL **16·2**

16.2.1 Self-Replicating Code 16·2
16.2.2 Actors: Origin of Malicious Code Threats 16·2
16.2.3 Actors: Structured Threats 16·3
16.2.4 Actors: Unstructured Threats 16·3
16.2.5 Access versus Action: Vector versus Payload 16·3

16.3 SURVEY OF MALICIOUS CODE **16·4**

16.3.1 Viruses 16·4
16.3.2 Worms 16·7
16.3.3 Trojans 16·8
16.3.4 Spyware 16·9
16.3.5 Rootkits 16·10
16.3.6 IRC Bots 16·10
16.3.7 Malicious Mobile Code 16·11

16.4 DETECTION OF MALICIOUS CODE **16·11**

16.4.1 Signature-Based Malicious Code Detection 16·11
16.4.2 Network-Based Malicious Code Detection 16·11
16.4.3 Behavioral Malicious Code Detection 16·12
16.4.4 Heuristic Malicious Code Detection 16·12

16.5 PREVENTION OF MALICIOUS CODE ATTACKS **16·12**

16.5.1 Defense in Depth 16·12
16.5.2 Operational Controls for Malicious Code 16·12
16.5.3 Human Controls for Malicious Code 16·13
16.5.4 Technical Controls for Malicious Code 16·13

16.6 CONCLUSION **16·14**

16.7 FURTHER READING **16·14**

16.8 NOTES **16·15**

16.1 INTRODUCTION. Malicious logic (or code) is "hardware, software, or firmware that is intentionally included in a system for an unauthorized purpose."[1] In this chapter, we enumerate the common types of malicious code, sources of malicious code, methods of malicious code replication, and methods of malicious code detection.

A 2011 study of 200 small and medium businesses (SMBs) with up to 249 employees reported that "… two in five SMBs know with certainty that they have suffered some sort of security breach as a result of employees navigating to Web sites that host malware, infected downloads or have been corrupted by malicious code."[2]

Common types of malicious code include viruses, worms, Trojan horses, spyware, rootkits, and bots. Emerging malicious code threats include kleptographic code, cryptoviruses, and hardware-based rootkits. Present-day malicious code threats do not always fit into neat categories, resulting in confusion when discussing the topic. It is not possible to classify all code as being *good* code or *malicious* code. Absent the *mens rea*, or criminal intent of the author or user, code is neither good nor bad. Authors develop code to achieve some goal or fulfill some purpose just as users run code to achieve some goal or purpose. It is therefore the context of use and the intent of the wielder that determines whether code is malicious.

16.2 MALICIOUS CODE THREAT MODEL. In a threat model or profile, an *actor* uses *access* to target an *asset*, with an *action*, to yield an *outcome*.[3] To understand the scope of the problem (and possible prevention), it is useful to study malicious code threats in this model. *Actors* may be structured or unstructured threats posed by individuals, organizations, or nation-states. *Access* is some allowed physical or logical path to the targeted *asset*. The execution of malicious code or logic is an action used to yield the desired *outcome*. This outcome could be intelligence (such as theft of trade secrets), surveillance, reconnaissance, disruption of operations, destruction of assets, publicity for some cause, or negative publicity for the victim.

16.2.1 Self-Replicating Code. Self-replicating code is not inherently malicious. A good deal of artificial intelligence research focuses on iterative self-replication. Hewlett-Packard and others[4] have researched how techniques of self-replication could yield beneficial software such as code-patching worms. John Von Neumann proposed the basic concept of self-replicating code in his 1949 paper, "Theory of Self-Reproducing Automata."[5] In 1961 at Bell Labs, Douglas McIlroy, Victor Vyssotsky, and Robert H. Morris played a game called Darwin in which two opposing programs (memory worms) would enter a system and only one would leave. In an interview with one of the authors, Robert H. Morris the former Chief Scientist of the National Security Agency (NSA), stated:

> We had this notion of putting two programs in a machine that would fight with each other and one would win and the other would die. Basically, the notion and the program were written by McIlroy and Vyssotsky and I was just on the side. But then, a few days later, I had a good idea and simply won the game and everyone else gave up … just because I happened to hit upon a very good idea for writing it.[6]

Although it would be desirable to live in a world free of war, the ever-increasing integration of technology and warfare necessitates the development of offensive information warfare capabilities. Malicious code implants can serve useful intelligence, surveillance, and reconnaissance capabilities in national security and law enforcement operations. In addition, by researching and developing malicious code techniques, practitioners can better prepare for defense against developing threats.

16.2.2 Actors: Origin of Malicious Code Threats. Prior to discussing the specific types of malicious code, it is worth understanding the origin and source of malicious code attacks. Malicious code may originate from structured or unstructured threats. Structured threats include nation-states, corporate criminals, and organized crime. Unstructured threats include rogue actors such as individual intruders and so-called script kiddies.

16.2.3 Actors: Structured Threats. Structured threats are organized, well funded, and may operate with a long-term strategic view. Structured malicious code threats tend to have intelligence, surveillance, and reconnaissance capabilities among their primary functions. Structured threats may use these capabilities to engage in industrial espionage, adversarial information operations, or serious ongoing fraud and theft.

Organized crime is responsible for 90 percent of malicious code threats.[7] Extortionists target the online gambling industry with threats of distributed denial of service (DDoS) attacks launched from compromised personal computers. Criminals use compromised systems to engage in pump-and-dump stock fraud with one media outlet reporting annual gains by such groups approaching $1 billion a year.[8] Malicious software-for-hire incidents are on the rise, indicating a maturing marketplace for malicious code. In 2006, an Israeli court sentenced a couple to prison time and ordered them to pay fines for authoring and placing malicious code used to spy on corporations and individuals by members of telecommunications and trading firms.[9]

Coordinated, systematic attacks originating from China routinely target defense and research and development facilities. In the 1999 text *Unrestricted Warfare*, two Chinese air force officers called for a transformation of warfare that would involve ongoing technological attacks on western assets.[10] A Chinese military white paper released in 2006 called for a strategy whereby China could win an "informationized war"[11] against the West that is "high-paced, high-technology, and digitized." Although Chinese officials publicly deny sponsoring such attacks, the lack of law enforcement action indicates, at the very least, toleration for electronic attacks on Western assets.

16.2.4 Actors: Unstructured Threats. Unstructured threats include rogue actors not acting in concert or coordination with larger entities. Although serious attacks may result from unstructured threats, they do not pose the same long-term challenges as structured threats. In testimony before the United States Congress, former NSA director Kenneth Minihan stated:

> The unstructured threat is random and relatively limited. It consists of adversaries with limited funds and organization and short-term goals. While it poses a threat to system operations, national security is not targeted. This is the most obvious threat today. The structured threat is considerably more methodical and well-supported. While the unstructured threat is the most obvious threat today, for national security purposes we are concerned primarily with the structured threat, since that poses the most significant risk.[12]

16.2.5 Access versus Action: Vector versus Payload. Malicious code attacks involve a vector and a payload. In biology, a vector is an agent that transfers (potentially harmful) material (code) from one location to another. In computer attacks, a vector is an avenue of *access*, such as an allowed path via physical access or via the network. Physical access occurs via internal personnel or others with access to the premises. Access via the network may occur via an allowed path to a Web server, an allowed path from a malicious Web server to a Web client, to a user via an email attachment, or to some other software process through an accessible port. The payload is a function (*action*) placed on the system to achieve some end. Payloads may include additional malicious logic, remote access software (rootkits and the like), or remote control (robot or bot) software to achieve some objective (spamming, DDoS attacks, and so forth).

16.3 SURVEY OF MALICIOUS CODE

16.3.1 Viruses. In biology, a *virus* is "[a]n infectious organism that is usually submicroscopic, can multiply only inside certain living host cells, and is now understood to be a non-cellular structure lacking any intrinsic metabolism and usually comprising a DNA or RNA core inside a protein coat."[13] In computer terms, a virus is self-replicating code that requires a host executable or document and the aid of a human to replicate. Joe Dellinger created the first virus for the Apple disk operating system in 1982 while a student at Texas A&M.[14] Fred Cohen created the first VAX computer virus in 1983 as a part of his doctoral research.[15] Len Adleman, the A in RSA, first applied the biologic metaphor virus to describe the results.

Types of viruses include boot sector, file infector, macro virus, logic bomb, cross-site scripting viruses (really a form of worm), polymorphic viruses, and cryptoviruses. However, modern malicious code threats tend not to fall so neatly into categories.

16.3.1.1 Boot Sector Viruses. When users boot a computer from digital media, they allow code present in the boot sector of the media to run on the microprocessor with very little intermediation. Boot sector viruses in the 1980s and 1990s used this mechanism to spread via infected removable media such as (now extinct) floppy disks. If a user errantly booted from an infected floppy disk, compact disc (CD), DVD, or flash drive, the virus would copy itself to the boot sector of the hard drive. Thereafter, the malicious program code would copy itself to each medium inserted into the system. Although reportedly still in existence, boot sector viruses comprise very few modern malicious code threats. *Virus Bulletin* reports, "The last boot sector viruses fell off the WildList in early 2006, but various types continue to appear in our prevalence reports and a batch of laptops infected with 'Stoned.Angelina' was released in Germany and Denmark in mid-2007. More recently, [T]rojans have been observed using similar techniques to plant rootkits to hide their activities."[16]

16.3.1.2 File Infector Viruses. File infector viruses inserted themselves into programs present on the host system. Whenever a user ran the host program (usually an .EXE or .COM file), the malicious code used the opportunity to insert itself into random access memory (RAM) and then replicate to other files on the system. Although probably still in existence, file infector viruses comprise relatively few modern malicious code threats.

16.3.1.3 Macro Viruses. Macro viruses spread via the macro definition languages used by some applications. The most widely abused is the Visual Basic for Applications (VBA) scripting language that Microsoft developed to allow the automation of functions inside the Office product suite. Any feature that allows for automation is a likely point of attack. Developers should carefully weigh security and ease of use when including such features in products. The first macro virus to target Microsoft Word[17] appeared in the wild in 1995. Since that time, macro viruses have comprised a significant number of successful attacks.

16.3.1.4 Logic Bombs. A logic bomb is a form of malicious code function sometimes built into viruses that wait for some sequence of events to activate such as disappearance of an employee record from the human resources database or a particular

date and time. Kabay gave some examples of early logic bombs in a 2002 Network World article:

> In 1985, a disgruntled computer security officer at an insurance brokerage firm in Texas set up a complex series of Job Control Language (JCL) and RPG (an old programming language) programs described later as "tripwires and time bombs." For example, a routine data retrieval function was modified to cause the IBM System/38 midrange computer to power down. Another routine was programmed to erase random sections of main memory, change its own name, and reset itself to execute a month later.
>
> In 1992, a computer programmer was fined $5,000 for leaving a logic bomb at General Dynamics. His intention was to return after his program had erased critical data and get paid lots of money to fix the problem.
>
> Time bombs are a subclass of logic bombs that "explode" at a certain time. Some of the first viruses, written in the 1980s, were time bombs. For example, the infamous "Friday the 13th" virus was a time bomb; it duplicated itself every Friday and on the 13th of the month, causing system slowdown. In addition, on every Friday the 13th it also corrupted all available disks. The Michelangelo virus from the early 1990s—one of the first viruses to make it into public consciousness because of news coverage—tried to damage hard disk directories on the 6th of March. The Win32.Kriz.3862 virus, discovered in 1999, detonates on Christmas day; its payload includes massive overwriting of data on all data storage units and also damage to the BIOS.
>
> In 2000, a Stamford, Conn., man was indicted in New York State Supreme Court in Manhattan on charges of unauthorized modifications to a computer system and grand larceny. The defendant worked for Deutsche Morgan Grenfell starting in 1996 as a programmer. By the end of 1996, he became a securities trader. The indictment charged that he inserted a programmatic time bomb into a risk model on which he worked as a programmer; the trigger date was July 2000. The unauthorized code was discovered by other programmers, who apparently had to spend months repairing the program because of the unauthorized changes the defendant allegedly inserted.[18]

In 2002 an employee of UBS PaineWebber planted a logic bomb in his employer's servers as part of an attempted stock manipulation scheme.[19] The perpetrator, Roger Duronio, purchased numerous put option stock contracts allowing him to sell UBS stock at a fixed, high price. On March 4, 2002, at 9:30 A.M., the logic bomb began deleting files on over 1,000 computers, causing a reported $3 million in damage. Duronio thought that the effects of the logic bomb would bring down the stock value of UBS, allowing him to make a large profit from his stock options. Duronio's plot failed and his profit did not materialize. Federal authorities later charged him with securities and computer fraud. In 2006 a judge sentenced Duronio to 97 months in prison for the attack.[20]

In March 2013, a logic bomb overwrote data on hard drives in South Korean banks and broadcasters. Kim Zetter of WIRED wrote,

> The logic bomb dictated the date and time the malware would begin erasing data from machines to coordinate the destruction across multiple victims, according to Richard Henderson, a threat researcher for FortiGuard Labs based in Vancouver, the research division of the security firm Fortinet.
>
> The attack, which struck machines on March 20, wiped the hard drives and master boot record of at least three banks and two media companies simultaneously. The attacks reportedly put some ATMs out of operation, preventing South Koreans from withdrawing cash from them.
>
> The malware consisted of four files, including one called AgentBase.exe that triggered the wiping. Contained within that file was a hex string (4DAD4678) indicating the date and time the attack was to begin—March 20, 2013 at 2 pm local time (2013-3-20 14:00:00). As soon as the internal clock on the machine hit 14:00:01, the wiper was triggered to overwrite the hard drive and master boot record on Microsoft Windows machines and then reboot the system.[21]

16.3.1.5 Cross-Site Scripting Viruses (or Worms). Cross-site scripting (XSS) viruses (or worms) replicate via flawed Web application servers and client code. A cross-site scripting exploit involving the social networking site MySpace and flaws in Microsoft Internet Explorer occurred in 2005 whereby a user named Samy amassed over 1 million friends overnight using a JavaScript insertion bug.[22] The Los Angeles Superior Court sentenced the author, Samy Kamkar, to three years' probation and 90 days of community service for his actions.

In an excellent review of XSS, Sherif Koussa wrote,

> Cross-site scripting is an attack that targets the application users. The simplicity of cross-site scripting is also why it is so powerful. If the application is vulnerable to cross-site scripting, the developer is not in charge of what runs on the user's browser anymore … the attacker is. Cross-site scripting could be used in attacks like authentication hijacking and session hijacking. The power of cross-site scripting manifests itself even more when combined with cross-site request forgery. …[23]

He then cites the case of the Twitter attack below:

> In 2009, Mikey Mooney, then 17 years old, claimed responsibility for attacking Twitter using XSS techniques: "… at least four separate variants of the original StalkDaily.com XSS worm hit the popular micro-blogging site Twitter, automatically hijacking accounts and advertising the author's web site by posting tweets on behalf of the account holders, by exploiting cross site scripting flaws at the site." Writer Dancho Danchev provided more details about the attack in his ZDNet article.[24]

16.3.1.6 Polymorphic Viruses. Polymorphic code modifies itself to evade detection. The virus code may accomplish this by dynamically reassembling itself to modify the underlying structure while retaining overall functionality. In other methods, the virus is encrypted, encoded, or packed, and a stub loader decrypts, decodes, or unpacks the virus at run time. UPX, the Ultimate Packer for eXecutables, is currently the most widely used packer format.

In the SOPHOS Security Threat Report 2013, the authors write:

> Polymorphism is not a new idea—malware authors have been using it for 20 years. Simply stated, polymorphic code changes its appearance in an attempt to avoid detection, without changing its behavior or goals. If a program looks different enough, attackers hope, antivirus software might miss it. Or the antivirus software might be forced to generate too many false positives, leading users to disable it.
>
> In a polymorphic attack, code is typically encrypted to appear meaningless and paired with a decryptor that translates it back into a form that can be executed. Each time it's decrypted, a mutation engine changes its syntax, semantics, or both.
>
> For instance, Windows malware authors have often used structured exception handling to obfuscate control flow and make it tougher to perform static analysis of programs before they run.
>
> Traditional polymorphic viruses are self-contained and must contain the mutation engine in order to replicate. Sophos and other security companies have become adept at detecting these forms of malware. With access to the mutation engine, it's easier to analyze its behavior.
>
> Today attackers are rapidly moving to web-distributed malware relying on server-side polymorphism (SSP). Now, the mutation engine and associated tools are hosted entirely on the server. Criminals can use these tools to create diverse file content on the fly. Recipients of this content (whether it is a Windows .exe, Adobe PDF, JavaScript, or anything else) see only one example of what the engine can create. They don't get to see the engine itself.[25]

16.3.1.7 Cryptographic Viruses. Cryptoviruses use encryption to encrypt data, making it inaccessible to the user. Although some viruses use symmetric algorithms for self-encryption to evade detection, such algorithms are inappropriate for data encrypting viruses, as any copy of the virus is going to yield a copy of the symmetric key. For this reason, proper cryptoviruses use asymmetric techniques. The Gpcode virus encrypts files by extension (.xls, .doc, and the like) and leaves behind a text message prompting the user to email a given address to pay a ransom for access to their files. The next generation of cryptoviruses will likely use hybrid cryptosystems.

In a 2012 report on cryptoviruses, researchers developed proof-of-concept models of viruses using the public-key cryptosystem (PKC) to conceal themselves from scanners and then activate in response to an encrypted key or ticket. Their list of applications of cryptography in malware includes:

- Resist reverse engineering
- Improve anonymity of communications from controllers to the malware
- More effective data theft and denial-of-service attacks
- Remote-control back doors for extortion[26]

16.3.2 Worms. The term *worm* refers to any form of self-replicating code that does not integrate into executable code. Common worm vectors include vulnerable services, email, instant messaging applications, and open file shares. Many worms use multiple vectors. For example, the Nimda worm spread via email, Web server vulnerabilities, hosted malicious Webpages, and open file shares.[27]

The first large-scale Internet worm infection was the Morris (aka Internet) Worm released by Robert T. Morris on November 2, 1988, while he was a student at Cornell University. The worm exploited flaws in the Sendmail and finger services to replicate and infected nearly 10 percent of Internet hosts. Although he claimed that this was an experiment gone awry, Morris became the first person convicted under the Computer Fraud and Abuse Act.

In 2001, Nicholas Weaver coined the term "Warhol Worms"[28] for those worms that had the capability of propagating very quickly (taking a cue from Warhol's famous quip that in the future everyone would have 15 minutes of fame). The SQL Slammer worm became the first such worm when it infected approximately 90 percent of vulnerable systems in 10 minutes.[29] The Slammer worm replicated due to a flaw in the Microsoft SQL database server. Because this was installed along with other components like Visual Studio, many people did not even know that they were running an SQL server. Slammer was exceedingly effective, because a Microsoft service pack downgraded a previously patched dynamic link library (dll) to an older, vulnerable version. Specifically, Microsoft issued four updates to *ssnetlib.dll* in 2002. Unfortunately, in October, hotfix Q317748 downgraded *ssnetlib.dll* to a vulnerable version, ironically making those who most faithfully applied patches and hotfixes the most vulnerable to Slammer.[30]

The SQL Slammer worm is an interesting case study because it was a vector without a payload. The 376-byte worm ran in memory without touching the hard disk, leading some to believe that Slammer was an experiment in propagation techniques. Since it was based on the User Datagram Protocol (UDP) and traveled in a single packet, the overhead was minimal and the propagation rate was greater than any previous malicious code threat.[31] The fact that the author(s) released it on a Saturday is also curious. It is unknown why a malicious attacker would deliberately release such a rabid

virus on a day of the week that would lessen the overall business impact. If the author had released the worm on a peak business day such as a Tuesday, the damage caused would have been much more severe.

Worms are used to distribute other forms of malware such as Trojans, spyware, rootkits, and remote command and control channels called *bots*. An indication of the overall maturity of malicious code attacks is the Bagel worm, which has been in production at least since 2004. Bagel, like many other mail worms, arrives as an email with a malicious attachment. When the user runs the attachment, the payload is executed, which does a number of things depending on the variant. Later variants of Bagel install an open application framework that the remote attackers may update and extend. The harvested systems deliver spam (unsolicited commercial email), harvest additional addresses, and act as a staging point for other attacks. The author(s) appear to follow a sophisticated software development methodology and testing process.[32] In 2007, attackers released 30,000 new variants in a six-week time span.[33] In April 2012, the Flashback worm "infected more than 650,000 Mac OS X systems using a vulnerability in Apple's version of Java."[34] Analysts with F-Secure wrote,

> Flashback is the most advanced OS X malware we've ever seen. It boasts a series of firsts for its kind. It was both the first to be VMware-aware and the first to disable XProtect, OS X's built-in malware protection program. Both these features were removed from later variants (the former presumably to avoid heuristic detections, and the latter presumably once the authors realized it was unnecessary, as XProtect was not designed to protect against non-quarantine fi les). Their removal indicates that Flashback is actively being reviewed and improved by its authors.
>
> Another interesting first is Flashback's exploitation of an unpatched vulnerability in the Java distribution of OS X, which allowed it to infect more than 650,000 Macs around the world. … This made Flashback roughly as common for Macs as Conficker was for Windows. … This means Flashback is not only the most advanced, but also the most successful OS X malware we've seen so far.
>
> Flashback's infection strategy is explicitly designed to select unprotected systems and will not infect a machine if certain security software or analysis tools are found. This implies that Flashback's authors are targeting less security-conscious users, at the expense of the total number of potential targets. This turns out to be an effective strategy, as security researchers had difficulties getting sufficient samples from users. It took a mistake on the part of Flashback's author to alert users to the presence of an infection and subsequently, to lead to the mass discovery of the malware.[35]

16.3.3 Trojans. A Trojan horse application, like the horse of Greek mythology, carries both an overt function and a covert function. Although attackers may use worms as one possible propagation vector, Trojans require that a user run the malicious program in order to be effective. Trojans tend to use some form of social engineering or manipulation to persuade the user to run the program. Email worms may appear to originate from a known associate and thereby trick the user. Other Trojans may take the form of games, free offers, pictures of popular celebrities, or files on peer-to-peer file sharing services. One study of the *Limewire* peer-to-peer file sharing service found that "68% of all downloadable responses containing executable, archival, and Microsoft Office file extensions"[36] contained malware. Queries for movies were most likely to hold malicious code.

The covert function of Trojan horse application is typically some form of a remote access Trojan, keylogger, dialer, IRC bot, or rootkit. Remote access Trojans provide full remote access to the system. The developers of the Bo2 K Trojan bill it as "the

most powerful network administration tool available for the Microsoft environment" and insist upon the fact that it is functionally no different from other remote access solutions.[37] Keylogging Trojans record keystrokes and periodically upload the data to a remote user. Attackers carried out the Windows 2000 source code theft using credentials stolen by a QAZ Trojan installed on a remote workers computer. A dialer is a form of Trojan that silently dials remote toll numbers and runs up a large telephone bill for the victim. Internet relay chat (IRC) bots act as autonomous IRC clients. Early IRC bots provided technical support and channel monitoring features but are now widely used for malicious purposes such as command and control capabilities.

In 2013, Kaspersky Labs reported on "a Trojan build specifically for Android smartphones."[38] sean Gallagher wrote,

> On March 25, the e-mail account of a Tibetan activist was hacked and then used to distribute Android malware to the activist's contact list. The e-mail's lure was a statement on the recent conference organized by the World Uyghur Congress. … If the targets opened the attachment … they received malware packaged in an Android APK file.
>
> When opened, the Trojan installs an app called "Conference" on the Android devices' desktops. If the app is launched, it displays a fake message from the chairman of the WUC—while sending back a message to a command and control server to report its successful installation. The malware provides a backdoor to the device via SMS messages sent by the server. On command, it returns the phone's contact lists, call logs, data about the smartphone, its geolocation data, and any SMS messages stored on it to a server via a Web POST upload.
>
> The server itself is running on a Chinese-language configured Windows Server 2003 machine sitting in a data center in Los Angeles. In addition to providing an upload point for the data stolen from Android devices, it also hosts more Android malware in its home page and provides a public Web interface (in Chinese) that allows direct control over phones that have been infected with the malware. While the server itself is at an IP address registered to a company called Emagine Concept, a domain pointed at the machine is registered to Shanghai Meicheng Technology Information Development Co., Ltd., a Chinese company with a contact in Beijing.[39]

This Trojan thus distributes spyware, the next topic.

16.3.4 Spyware. The term *spyware* refers to any software that collects user information without consent. Common varieties of spyware collect information on Web usage, serve advertising content (pop-ups), log keystrokes, engage in click fraud, or monitor program usage and licensing. Unauthorized access or exceeding authority on a computer system is a violation of the Computer Fraud and Abuse Act. Although some spyware may be illegal, some developers insist that they are engaged in a legitimate business activity. Such spyware provides an end user license agreement (EULA) allowing the user to opt out of installation. Since the user must grant permission for the software to install, it is almost certainly legal. However, whether this is ethical and conscionable is a different matter. Sony attracted significant negative publicity for their use of spyware technology to limit the ability of listeners to copy music compact discs. Due to this activity, Sony faced legal charges in multiple states and in 2005 settled a class-action lawsuit. Informed consent is the rule of thumb in any monitoring system. Organizations should endeavor to be forthright in these matters in order to maintain a positive public profile and to avoid legal challenges.

In May 2013, news surfaced about the use of spyware to monitor the communications of a Bahrain human-rights activist. In a court hearing, details arose describing

how Dr Ala'a Shehabi was monitored through infection of her systems through a phishing attack:

> According to her witness statement, a few weeks after her arrest Shehabi received a series of emails, the first purportedly from Kahil Marzou who was the deputy head of Bahrain's main opposition party, including one containing a virus. Other emails that claimed to be from an Al Jazeera journalist were also infected. Research found that the emails contained a product called FinSpy, distributed by a British company, Gamma International.
>
> The witness statement claims that when a person's computer is infected with FinSpy, "it allows access to emails, social media messaging, and Skype calls, as well as copying the files saved on the hard disk. These products also enable whoever is doing the targeting to commandeer and remotely operate microphones and cameras on computers and mobile phones."[40]

16.3.5 Rootkits. A rootkit consists of a set of tools for covertly compromising a system and maintaining administrative (root) access for the intruder. Present-day rootkits compromise a system at the application, library, kernel, hypervisor, or hardware level. Early application-level rootkits targeted UNIX-like systems and replaced standard system utilities (netstat, ps, and the like) with versions that omitted information on the intruder, such as open ports, running processes, open files, and other activity. API-level rootkits modify or patch the system call table to redirect system calls (like an API-level monkey-in-the-middle attack). Kernel-level rootkits run as device drivers and dynamically load into the kernel; compromising the integrity of the core of the operating system to filter information presented to users via application-level processes.

Nearly all modern operating systems make use of virtualization for memory and process management. An operating system normally runs in ring 0 of the microprocessor; the most privileged level. A hypervisor is a layer of code between the operating system and the hardware that fools the operating system into believing that it is running in ring 0. The hypervisor monitors and arbitrates exchanges between virtual machines and the real hardware. Hypervisor design is a key part of the trusted computing framework, and microprocessor designers like Intel and AMD are including enhanced hardware-based virtualization capabilities in their products. An emerging generation of rootkit technology uses this framework to create potentially undetectable code.[41,42]

Other potential rootkit threats come from insiders, hardware designers, and manufacturers. An attacker with physical access could insert a peripheral component into a computer to create a logically undetectable rootkit. A hardware manufacturer or designer could include rootkit-like capabilities in a microcircuit. The 2005 sale of IBM's Personal Computer Division to Chinese manufacturer Lenovo caused enough concern in the United States to prompt a national security review of the transaction.[43] The review by the House Committee on Foreign Investment in the United States (CFIUS) eventually approved the sale despite concerns expressed by Reps. Henry Hyde and Don Manzullo of Illinois and Rep. Duncan Hunter of California.[44]

16.3.6 IRC Bots. IRC bots are autonomous agents that make use of Internet Relay Chat (IRC) to offer interactive services like channel monitoring, support, information services, and games. Greg Lindhal wrote the first IRC bot GM (game master), which led users through the text-based role-playing game "Hunt the Wumpus." In 1999, the PrettyPark worm became the first worm to use IRC as a remote control channel. Infected systems would check into an IRC server and channel to download updates and upload stolen data. Attackers create malicious bots to carry out DDoS attacks (DDoS

bots), send unsolicited commercial email (SpamBots), and engage in exploitation, theft, or fraud.

Examples of current bot frameworks are GTbot, SDbot, Agobot, Goabot, Randex, Spybot, and Phatbot. The current generation of bot technologies includes keylogging, port scanning, exploitation, packet sniffing, process hiding, and adware fraud capabilities. A single network of these compromised systems may reach more than 100,000 bots. Researchers in 2004 estimated that there were over one million bots currently connected to the Internet.[45] Bot herders rent out these massive distributed networks systems to criminal organizations. In 2006, a California court sentenced one such bot herder, Jeanson Ancheta, to five years in prison for conspiring to violate the Computer Fraud and Abuse Act, conspiring to violate the CAN-SPAM Act, causing damage to computers used by the federal government in national defense, and accessing protected computers without authorization, in order to commit fraud.[46]

16.3.7 Malicious Mobile Code. Web servers may host pages containing malicious mobile code. ActiveX controls, Java applets, JavaScript, Adobe Flash animations, and any other type of dynamic executing code can download to a user's system and run in their context with all associated privileges. Malicious Web servers are one method for dropping Trojans and bots onto computers. An attacker may use a spam or phishing attack to encourage users to click on a link embedded within an email. Another vector is to use a similar domain name to a legitimate organization. Yet another vector is for the site to offer information in some matter. Once indexed in the major search engines, this content will draw a number of users to the site. See Chapter 17 in this *Handbook* for further details on malicious mobile code.

16.4 DETECTION OF MALICIOUS CODE. Common methods of detecting malicious code include signature-based, network-based, and behavioral heuristic techniques. However, as far back as in the 1984 work *Computer Viruses—Theory and Experiments*, Fred Cohen demonstrated that the only way one could prevent all possible viral code would be through isolation.[47] While there are numerous methods of detecting malicious code, no one technique works on every variety or in every circumstance, and there is always some method for successfully evading detection.

16.4.1 Signature-Based Malicious Code Detection. Some of the oldest methods of malicious code detection are signature-based methods that utilize known strings or patterns in the code. Signature-based methods are easy to implement and impose very low overhead on the system but are just as easily evaded. Polymorphism and metamorphism[48] are two methods by which malware may change form over time and thereby evade signature-based detection. Detection systems can use hash functions to fingerprint a given binary program or code fragment. However, the hash will fail to match on any modified version, making this method reliable but not always useful. Overall, signature-based methods are not terribly reliable, although signatures are one useful metric in more complex heuristics.

16.4.2 Network-Based Malicious Code Detection. Network-based methods of malicious code detection look for network artifacts associated with malicious code, such as a connection to a server as in the case of a keylogging Trojan or Internet Relay Chat (IRC) bot. Network anomaly detection works well but is expensive and poorly understood by many security practitioners. A simple method for

network-based detection is to analyze network flow (netflow) data in comparison with a statistical database of known good traffic. However, malicious code that acts normal can bypass such methods of detection. Predicting what is normal for a given environment may be difficult for an outsider, but analyzing certain activity like domain name system, email, and Web usage can produce generic models for evasion.

16.4.3 Behavioral Malicious Code Detection. Behavioral methods of detecting malicious code analyze the actions of running software to look for illegitimate activity. This could be opening a port, connecting to a remote host, or modifying the system call table or other memory areas. Behavioral methods of detection will fail if the malware acts normal or if the malicious code can target and exploit the detection system itself. A slow but effective method of detecting potentially malicious code is to use a virtual machine approach, whereby the system allows code to execute in a sandboxed virtual machine. This allows for a full functional analysis of the code, but the current throughput of such systems makes them not useful for today's high-speed production environments.

16.4.4 Heuristic Malicious Code Detection. Heuristics that are more complex may use both statistical and behavioral models to determine a relative score to a normal corpus (statistical database in this case) of legitimate behavior. Bayesian analysis is widely used in the detection of spam. However, this method can also detect new variants of existing malicious code with a high degree of accuracy. N-gram analysis is a form of frequency analysis borrowed from natural language processing that can model software and fingerprint data types. This method is useful for detecting executable code embedded within other data objects. In short, there is no single, monolithic way to detect all malicious code. Spinellis demonstrated that reliable detection of malicious code is NP-complete,[49] meaning that the dilemma is not solvable in polynomial time. However, the technologies discussed, used in concert, are relatively effective in detecting many common malicious code threats.

16.5 PREVENTION OF MALICIOUS CODE ATTACKS

16.5.1 Defense in Depth. As the problem of malicious code is demonstrably NP-complete (not solvable), a single antivirus program cannot protect against all malicious code threats. One strategy, called *defense in depth*, uses operational, human, and technical controls. Another strategy is to build networks and applications that only function one correct way. Such *orthogonal* networks and applications are a rarity as they are expensive and difficult to design, and many firms will resist imposing rigid limitations on the enterprise.

16.5.2 Operational Controls for Malicious Code. All organizations must create written policies and procedures regarding the introduction of program code into the operating environment. Policies should define what persons the firm allows to install programs, acceptable use for Internet access, acceptable use for email systems, and what to do if users suspect the compromise of a system or user. Organizations should subject all new employees to some level of background investigation. This should include, at a minimum, a criminal records search, verification of all references and credentials, and a credit report. Employers should ensure that the prospective employees are forthright and truthful on their application. If employees lie prior to employment, the odds are that they will lie later as well.

16.5.3 Human Controls for Malicious Code. All users (including executives) should receive training on the policies and procedures of the organization. Due to the evolving nature of the malicious code threat, the organization should update and refresh this training annually at the very least. The training sessions should introduce the prevalent types of threats, how to detect the threats, and proper response. Currently, this training should include identification of advance-fee fraud (also known as Nigerian 419 frauds), social engineering attempts, and detection of malicious attachments. Users should notify the help desk or other entity if they encounter any anomalous system or user behavior.

16.5.4 Technical Controls for Malicious Code. For more detailed information on antivirus technology, see Chapter 41 in this *Handbook*.

16.5.4.1 Implementing Antivirus Systems. Implement an antivirus (A/V) solution that suits the operational environment. This solution should include both network-based and host-based systems. Network-based systems function inline like a gateway, while host-based systems run on host end points. These systems should come from different vendors, as there is less benefit when using the same software on the network and hosts. A diverse detection and containment strategy will help firms avoid the pitfalls associated with active malcode exploitation of antivirus software, as occurred in 2004[50] and again in 2006.[51] The last thing that organizations should permit when designing an A/V strategy is for antivirus software to serve as the vector for an active malcode attack.

The A/V solution should provide a mechanism for dynamically applying updates and should do so daily, at the least. Email systems may require a separate inline appliance to detect email–borne malicious code as well as spam, fraud, and phishing attacks. By using a diversity of detection approaches, organizations can attain higher rates of detection than by using a single stand-alone product.

16.5.4.2 Host Configuration Controls and Security. Host configuration can mitigate many malware threats prior to emergence. Implement a form of automatic updates (patches, etc.) that supports the operating environment. Many technical malicious-code threats (e.g., worms) target well-known and patched flaws. Eliminate all noncritical software and services. This will help minimize threats that target the ever-burgeoning code complexity that security professionals face. In one study, programmers trained in the capability maturity model for secure software development continued to make 4.5 errors per 1,000 lines of code.[52] For an operating system like Microsoft Vista, which one estimate puts at 50 million lines of code,[53] extrapolating this statistic means that there are likely at least 225,000 code errors. Disabling or removing as much of this code as possible is a reasonable preventive control. If the environment permits, remove all Web browsers. If the environment does not allow this, lock down the browser configuration and use a secure Web proxy.

16.5.4.3 Network-Based Security Controls. A layered defense of routers, firewalls, proxies, and switched virtual local area networks (VLANs) can mitigate malicious code propagation. At the router, filter all inbound bogus network addresses (BOGONs), Request for Comment (RFC) 1918 addresses, and spoofed internal addresses per RFCs 2267[54] and 3704.[55] If the enterprise faces specific threats from certain nations, filter the network blocks allocated to those nations at the border router.

Use the current best practices in firewall configuration. (These change rapidly.) Use a secure, authenticated Web proxy, and force all clients to use the proxy. (No user should be able to access the Web directly.) Disable all unused LAN access ports, or consider using 802.1x authentication to do so automatically. Segment the network into functional workgroups using VLANs or physically separate switch architecture. If possible, implement access control lists between VLANs or switched segments. The goal is to make the network as orthogonal as possible. If the network can function only one right way, the organization will avoid many automated malicious code threats entirely.

16.5.4.4 Network Monitoring. Prevention of malicious code threats is ideal, but detection is critical. Most malicious code attacks will have some artifact, whether host based or network based. To detect these artifacts, organizations should establish a security information management system that aggregates device logs, server logs, host logs, intrusion detection system alerts, and network flow data. Network flow data is a useful tool for detecting anomalous network activity. A network flow is a 5-tuple consisting of a source address, destination address, source port, destination port, and protocol with an associated time stamp. Any malicious network activity is going to have an associated flow. Stealthy malware, however, will attempt to act normal in order to evade detection. Detecting anomalous activity requires that operators understand what is normal for the environment. Without a historical statistical database or extensive experience, this can be difficult. Network anomaly detection (NAD) uses a statistical modeling methodology[56] to use this data to detect statistical outliers. This is extraordinarily effective for detecting new malicious code threats as well as other forms of anomalous behavior. However, a lack of understanding appears to limit the adoption of NAD. This is one area where many organizations could improve their practice of information security controls.

16.6 CONCLUSION.

Malicious code threats are as numerous as the variety of nonmalicious code. Prevention of all malicious code is not possible as the problem is demonstrably NP-complete.[57] However, a strategy of defense in depth that uses operational, human, and technical controls can be relatively effective. Used properly, the current generation of technical controls available to organizations is effective in stopping a majority of malicious code threats. However, the trusted insider typically has the access needed to turn a threat into a reality of operational risk. Current trends in malicious code threats indicate a continued pattern of organized criminal involvement and international espionage that continues to target the weakest link in the security chain: the human being.

16.7 FURTHER READING

Blunden, Bill. *The Rootkit Arsenal: Escape and Evasion in the Dark Corners of the System*, 2nd ed. Jones & Bartlett Learning, 2012.

Butler, Jamie and Greg Hoglund. *Rootkits: Subverting the Windows Kernel*. Addison-Wesley, 2005.

Harley, David., et al. *Avien Malware Defense Guide for the Enterprise*. Burlington, MA: Syngress, 2007.

Harley, David. *OS X Exploits and Defense*. Burlington, MA: Syngress, 2008.

Sikorski, Michael. *Practical Malware Analysis: The Hands-On Guide to Dissecting Malicious Software*. No Starch Press, 2012.

Szor, Peter. *The Art of Computer Virus Research and Defense*. Cupertino, CA: Symantec Press, 2005.

Young, Adam, and Moti Yung. *Malicious Cryptography: Exposing Cryptovirology*. Hoboken, NJ: John Wiley & Sons, 2004.

16.8 NOTES

1. NCSC, "Glossary of Computer Security Terms," Version 1, 1988, NCSC-TG-004 in the Rainbow Series, www.fas.org/irp/nsa/rainbow/tg004.htm
2. GFI, "GFI Software Survey: 40% of SMBs Have Suffered a Security Breach Due to Unsafe Web Surfing," *EON*. October 12, 2011, http://eon.businesswire.com/news/eon/20111012005707/en/web-reputation/web-security-survey/small-business-web-monitoring-software
3. C. Alberts and A. Dorofee, "OCTAVE Threat Profiles," 2001, www.cert.org/archive/pdf/OCTAVEthreatProfiles.pdf
4. R. Lemos, "Good Worms Back on the Agenda," *Security Focus*, January 27, 2006, www.securityfocus.com/news/11373
5. John von Neumann, "Theory of Self-Reproducing Automata," *Part 1: Transcripts of Lectures Given at the University of Illinois, Dec. 1949*, ed. A. W. Burks (University of Illinois, 1966). Urbana, IL.
6. R. H. Morris, Recorded interview with Robert Guess, July 2005.
7. D. Llet, "Antivirus Firm Says Organized Crime Growing Online," *ZDnet News*, December 9, 2004, http://news.zdnet.com/2100-1009_22-5486201.html (URL inactive).
8. J. Hoskyn, "SEC Targets Pump and Dump Scammers," *IT Week*, March 9, 2007, www.itweek.co.uk/vnunet/news/2185164/sec-targets-pump-dump-stock (site discontinued).
9. SOPHOS, "Court hands hefty fine and jail sentence to Israeli spyware couple, reports Sophos," *SOPHOS | Press Releases*, March 26, 2006, www.sophos.com/en-us/press-office/press-releases/2006/03/israelspyduo.aspx (URL inactive).
10. Quio Liang and Wang Xiangsui, *Unrestricted Warfare* (Beijing: PLA Literature and Arts Publishing House, 1999), www.terrorism.com/documents/TRC Analysis/unrestricted.pdf (Access Permission required).
11. J. Rogin, "Cyber Officials: Chinese Hackers Attack 'Anything and Everything,"' February 13, 2007, www.fcw.com/article97658-02-13-07-Web&printLayout (URL inactive).
12. Kenneth A. Minihan, "Prepared Statement before the Senate Governmental Affairs Committee," June 14, 1998, www.senate.gov/~gov_affairs/62498minihan.htm (URL inactive).
13. *Oxford English Dictionary,* 2007.
14. J. Dellinger, Usenet post, December 2, 1987, http://yarchive.net/risks/early_virus.html
15. F. Cohen, "Experiments with Computer Viruses," 1984, www.all.net/books/virus/part5.html
16. Virus Bulletin, "Boot sector virus," *Virus Bulletin | Glossary*. 2013, www.virusbtn.com/resources/glossary/boot_sector_virus.xml
17. Microsoft, "What to Do If You Have a Macro Virus," October 4, 2002, http://support.microsoft.com/kb/181080/en-us

18. M. E. Kabay, "Logic Bombs, Part 1." *Network World | Security Newsletter*. August 21, 2002, www.networkworld.com/newsletters/sec/2002/01514405.html
19. U.S. Department of Justice, United States Attorney District of New Jersey, "Disgruntled UBS PaineWebber Employee Charged with Allegedly Unleashing 'Logic Bomb' on Company Computers," December 17, 2002.
20. U.S. Department of Justice Press Release, "Former UBS Computer Systems Manager Gets 97 Months for Unleashing 'Logic Bomb' on Company Network," 2006, www.usdoj.gov/usao/nj/press/files/pdffiles/duro1213rel.pdf (URL inactive).
21. Kim Zetter, "Logic Bomb Set Off South Korea Cyberattack," *Wired*, March 21, 2013, www.wired.com/threatlevel/2013/03/logic-bomb-south-korea-attack
22. E. Lai, "Teen Uses Worm to Boost Ratings on MySpace.com," *Computer-World,* October 17, 2005, www.computerworld.com/securitytopics/security/holes/story/0,10801,105484,00.html
23. Sheif Koussa, "Security Code Review Techniques: Cross-Site Scripting Edition," *Developer Connection | Jonathan Rozenblit*, March 5, 2013, http://blogs.msdn.com/b/cdndevs/archive/2013/03/05/security-code-review-techniques-cross-site-scripting-edition.aspx
24. Dancho Danchev, "Twitter Hit by Multiple Variants of XSS Worm," *ZDNET | Security*, April 14, 2009, www.zdnet.com/blog/security/twitter-hit-by-multiple-variants-of-xss-worm/3125
25. SOPHOS, "Security Threat Report 2013: New Platforms and Changing Threats," *SOPHOS*, February 1, 2013, www.sophos.com/en-us/medialibrary/PDFs/other/sophossecuritythreatreport2013.pdf
26. Shafiqul Abidin, Rajeev Kumar, and Varun Tiwari. "Review Report on Cryptovirology and Cryptography," *International Journal of Scientific & Engineering Research*, 11, 2012: 1–4.
27. CERT. "Advisory CA-2001-26 Nimda Worm." Original release date: September 18, 2001, Revised: September 25, 2001. Source: CERT/CC, www.cert.org/advisories/CA-2001-26.html
28. N.Weaver, "Warhol Worms: The Potential for Very Fast Internet Plagues," 2001, www.iwar.org.uk/comsec/resources/worms/warhol-worm.htm
29. D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver, "The Spread of the Sapphire/Slammer Worm," CAIDA, 2003, www.caida.org/publications/papers/2003/sapphire/sapphire.html
30. S. Berinato, "Patch and Pray," *CSO Online*. August 14, 2003, www.csoonline.com.au/index.php/id;1337625166;fp;8;fpid;5
31. T. Vogt, "Simulating and Optimising Worm Propagation Algorithms" September 29, 2003, http://web.lemuria.org/security/WormPropagation.pdf
32. "Year of the Beagle: The Beagle Worm History, Part III," *Infection Vectors* Website, February 2005, www.infectionvectors.com/vectors/year_of_the_beagle.htm
33. CommTouch, "Malware Outbreak Trend Report: Bagle/Beagle," March 6, 2007, www.commtouch.com/documents/Bagle-Worm_MOTR.pdf (URL inactive).
34. Brian Krebs, "Who Wrote the Flashback OS X Worm?" *Krebs on Security*, April 3, 2013, http://krebsonsecurity.com/2013/04/who-wrote-the-flashback-os-x-worm
35. Broderick Ian Aquilino, "Flashback OS X Malware," *F-Secure*, October 5, 2012. www.f-secure.com/weblog/archives/Aquilino-VB2012.pdf

36. A. Kalafut, A. Acharya, and M. Gupta, "A Study of Malware in PeertoPeer Networks," 2006, www.imconf.net/imc-2006/papers/p33-kalafut.pdf (site discontinued).
37. B02 k Website, "A Note on Product Legitimacy and Security," ND, http://bo2 k.sourceforge.net/docs/bo2 k_legitimacy.html
38. Costin Raiu, Kurt Baumgartner, and Denis, "Android Trojan Found in Targeted Attack," *SECURELIST*, March 26, 2013, www.securelist.com/en/blog/208194186/Android_Trojan_Found_in_Targeted_Attack
39. Sean Gallagher, "First targeted attack to use Android malware discovered: Kaspersky uncovers trojan spread by 'spear-phish' to Tibet activists." *ars technica*, March 26, 2013, http://arstechnica.com/security/2013/03/first-targeted-attack-to-use-android-malware-discovered/
40. Jamie Doward, "UK company's spyware 'used against Bahrain activist', court papers claim: Human rights groups hope email evidence can force review of export controls on surveillance equipment," *The Guardian*, May 11, 2013, www.guardian.co.uk/world/2013/may/12/uk-company-spyware-bahrain-claim
41. S. King and P. Chen, "SubVirt: Implementing Malware with Virtual Machines," 2006. www.eecs.umich.edu/virtual/papers/king06.pdf
42. J. Rutkowska, "Subverting the Vista Kernel for Fun and Profit," 2006, www.blackhat.com/presentations/bh-usa-06/BH-US-06-Rutkowska.pd
43. A. Wolfe, "U.S. To Review IBM-Lenovo Sale," *CMP*, 2005, www.crn.com/hardware/59100372
44. IBM, "Committee on Foreign Investment in U.S. completes review of Lenovo-IBM deal," *Information Week*, 2005, www.informationweek.com/news/hardware/showarticle.jhtml?articleID=162400445 (URL inactive).
45. P. Bacher, T. Holz, M. Kotter, and G. Wicherski, "Know Your Enemy: Tracking Botnets," *The Honeynet Project* Website, 2005, www.honeynet.org/papers/bots
46. U.S. Department of Justice, "'Botherder' Dealt Record Prison Sentence for Selling and Spreading Malicious Computer Code," 2006, www.cybercrime.gov/anchetaSent.htm (URL inactive).
47. F. Cohen, "Computer Viruses—Theory and Experiments," 1984, www.all.net/books/virus/index.html
48. M. Christodorescu, S. Jha, S. Seshia, D. Song, and R. Bryant, "Semantics-Aware Malware Detection," 2005, www.eecs.berkeley.edu/~sseshia/pubdir/oakland05.pdf
49. D. Spinellis, "Reliable Identification of Bounded-length Viruses is NP-complete," 2003, www.dmst.aueb.gr/dds/pubs/jrnl/2002-ieeetit-npvirus/html/npvirus.html
50. D. Fisher, "Fast-Moving Worm Crashes Computers," *eWeek* Website, March 22, 2004, www.eweek.com/article2/0,1895,1551998,00.asp
51. T. Gray, "New Worm Attacks Through Symantec Antivirus App," *Newsfactor* Website, December 18, 2006, www.newsfactor.com/news/Worm-Attacks-Through-Antivirus-Flaw/story.xhtml?story_i=d121000E3SRW1 (URL inactive).
52. W. Humphrey, "Three Dimensions of Process Improvement Part II: The Personal Process," 1998, www.stsc.hill.af.mil/crossTalk/frames.asp?uri=1998/03/dimensions.asp
53. S. Manes, "Dim Vista," *Forbes Magazine,* February 26, 2007, www.forbes.com/free_forbes/2007/0226/050.html (URL inactive).

54. P. Ferguson and D. Senie, "RFC 2267: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," *Internet Engineering Task Force*, 1998, www.ietf.org/rfc/rfc2267.txt

55. F. Baker and P. Savola, "RFC 3704: Ingress Filtering for Multihomed Networks," *Internet Engineering Task Force*, 2004, www.ietf.org/rfc/rfc3704.txt

56. National Security Agency, "Network Anomaly Detection Algorithm," n.d., www.nsa.gov/techtrans/techt00029.cfm (URL inactive).

57. B. Aditya Prakash, Lada Adamic, Theodore Iwashyna, Hanghang Tong, and Christos Faloutsos. "Fractional Immunization in Networks." *Carnegie Mellon University | Computer Science*, January 24, 2013, www.cs.cmu.edu/~badityap/papers/smartalloc-sdm13.pdf

CHAPTER 17

MOBILE CODE

Robert Gezelter

17.1 INTRODUCTION 17·1
- 17.1.1 Mobile Code from the World Wide Web 17·3
- 17.1.2 Motivations and Goals 17·5
- 17.1.3 Design and Implementation Errors 17·5

17.2 SIGNED CODE 17·6
- 17.2.1 Authenticode 17·6
- 17.2.2 Fundamental Limitations of Signed Code 17·6
- 17.2.3 Specific Problems with the ActiveX Security Model 17·7
- 17.2.4 Case Studies 17·8

17.3 RESTRICTED OPERATING ENVIRONMENTS 17·11
- 17.3.1 Java 17·11
- 17.3.2 Virtual Machines 17·12
- 17.3.3 Operating Environment Summary 17·12

17.4 DISCUSSION 17·12
- 17.4.1 Asymmetric, and Transitive or Derivative, Trust 17·12
- 17.4.2 Misappropriation and Subversion 17·13
- 17.4.3 Multidimensional Threat 17·13
- 17.4.4 Client Responsibilities 17·14
- 17.4.5 Server Responsibilities 17·15

17.5 SUMMARY 17·16

17.6 FURTHER READING 17·16

17.7 NOTES 17·17

17.1 INTRODUCTION. Computing has been transformed in the last twenty years. In these past two decades, access to the World Wide Web has gone from a rare desktop novelty to the ubiquitous; now the majority of people carry mobile devices with full wireless Internet access. The online world has blossomed over the past two decades. The term mobile device is no longer an ironic oxymoron. Twenty years ago, today's dumb phone was the size of a suitcase; today, a far more powerful smart phone easily fits in a child's palm.

Today's portable devices are far more powerful than the mainframes of yesteryear. Their capabilities present a security and information assurance challenge. The consumer-friendly model of easy-to-add enhancements to these devices complicates the security ecosystem.

The World Wide Web began with simple, static pages. Applications soon started to produce pages on demand. The evolution of Web-related technologies has grown to include significant reliance on JavaScript (also known as ECMAScript[1]) through the use of AJAX and other techniques, as well as the use of Java and other plug-ins.

The challenge is that such code is often delivered by outside agencies as it is needed (e.g., JavaScript elements are downloaded from the supplier or other Web server as pages are loaded; apps for mobile devices are installed on the fly at little or no nominal cost when desired by the owner of the device). Thus, there is often little or no opportunity to vet the behaviors of such code, or even to maintain a semblance of control. One version of the code may be well-behaved; the next version may enable a compromise of the device, whether inadvertent or deliberate. The Apple App Store used with iPhones®, iPods®, and iPads® has an approval and vetting process, but the process has not been totally without flaw.[2] As an example, it is not uncommon to reference external JavaScript elements using unencrypted hypertext transfer protocol (HTTP), rendering the Web page inherently vulnerable to masquerading, man-in-the-middle, domain name impersonation, and other attacks. If nothing else, the integrity of an externally supplied element is dependent on the integrity and security of every element in the network, a daunting reality.

At the most basic level, mobile code is a set of instructions that are delivered to a remote computing device for dynamic execution or installation on the device. Such prepackaged components are a boon to applications on desktops and mobile devices; however, the inherent power of the approach comes with significant hazards. The hazards associated mobile code stem from its inherent ability to do more than just display characters on the remote display.

It is this dynamic nature of mobile code that causes policy and implementation difficulties. A blanket prohibition on mobile code is secure, but that prohibition would prevent users of the dynamic Web from performing their tasks and deny users of mobile devices much of the device's utility. The tension among integrity, flexibility, and dynamism is at the heart of the issue.

The advent of mass-market appliances, from personal devices (e.g., smart phones, tablets) to out-of-the-box appliances (e.g., file servers) makes blanket prohibitions often infeasible. Such devices are delivered with prepackaged or downloaded software, which can then be enhanced via the downloading of entire new versions, or extensions known as *apps*. In both cases, the software elements developed by the manufacturer or third parties are prepackaged, with little or no possibility for customer organization testing. Although IT consumerization has benefits, one of the significant drawbacks is the loss of user and corporate visibility into what software is in use, and what access the software has to information.[3,4,5,6]

Smart phones and tablets are emblematic of IT consumerization. Without the enhancements provided by apps these devices provide a faint shadow of their true capability. However, a rogue or weakly defended app can be downloaded by an unsuspecting user, compromising sensitive data. Apps are supplied as prepackaged, binary installation files, typically delivered over wireless communications pathways (e.g., Wi-Fi or the cellular communications network).

Several definitions, as used by United States military forces but applicable to all, are useful in considering the content of this chapter:

Enclave. An information system environment that is end to end under the control of a single authority and has a uniform security policy, including personnel and physical security. Local and remote elements that access resources within an enclave must satisfy the policy of the enclave.

Mobile code. Software obtained from remote systems outside the enclave boundary, transferred across a network, and then downloaded and executed on a local system without explicit installation or execution by the recipient. Mobile code is a powerful software tool that enhances cross-platform capabilities, sharing of resources, and Web-based solutions. Its use

> is widespread and increasing in both commercial and government applications. … Mobile code, unfortunately, has the potential to severely degrade … operations if improperly used or controlled.
>
> **Malicious mobile code.** Mobile code software modules designed, employed, distributed, or activated with the intention of compromising the performance or security of information systems, increasing access to those systems, providing the unauthorized disclosure of information, corrupting information, denying service, or stealing resources.[7]

17.1.1 Mobile Code from the World Wide Web. In discussions of the World Wide Web, the phrase *mobile code* generally refers to executable code, other than HTML and related languages (e.g., Extensible Markup Language, or XML), delivered electronically over a communications network (e.g., Web, email) for execution on the client's computer. The most common packaging technologies for mobile code are apps, ActiveX, Java, and JavaScript.

Mobile code can directly perform covert functions on a client system, accessing information or altering the operation or persistent state of the system; it can create accidental or deliberate vulnerabilities that can be exploited at a later time.

The widespread use of mail clients that support executable email content (e.g., HTML email, with either embedded or referenced program code), has become a significant source of vulnerability.

Pop-ups can also be a source of vulnerability in many ways; in particular, they can access other WWW sites and pages and may give rise to log entries that can cause legal problems. The experience of Julie Amero, a Norwich, Connecticut, substitute teacher, is as troubling as it is cautionary. Ms. Amero was accused[8] and convicted of using a classroom computer to access inappropriate material. The worst charges were subsequently withdrawn when it became clear that the classroom computer at the center of the incident had been infected with malware days before the incident, the antivirus subscription had lapsed, and the testimony of willful access was demonstrated to be factually incorrect. However, Ms. Amero was subjected to an extended prosecution and pilloried in the press. Settling the case required Ms. Amero to plead guilty to a misdemeanor, surrender her teaching license, and pay a fine.[9,10,11]

Although malicious software such as viruses, worms, and Trojan horse programs written in compiled, interpreted, or scripting languages such as Visual Basic also might be considered mobile code, these pests are not generally labeled as such; this chapter deals only with *apps*, ActiveX controls, Java applets, and JavaScript programs. See Chapter 16 in this *Handbook* for details of other kinds of malicious code.

Today's increasingly dynamic and rich Internet applications rely on AJAX,[12] WebSockets,[13] and other technologies, which in turn rely on JavaScript and other mobile code technologies. These applications increase the scope of the threat while also making it more difficult to ban the use of the vulnerable technologies.

Although systems or applications crashes are the most spectacular problems with mobile code, they are not the most dangerous. Persistent, silent, and covert access or modification of client system data pose far more serious threats. For example, some Web sites have covertly obtained email addresses from users' browsers, resulting later in unwanted commercial email. Similar attacks have been made against mobile devices phonebooks and contact lists. As the motivations for such attacks have migrated from pranks to criminal and nation-state-sponsored espionage, the importance of undetected covert access to data has correspondingly increased.

In the past, antimalware tactics have relied heavily on widespread distribution of threats. The expected emergence of designer mobile code, specifically targeted to a

particular system or a small number of systems, is a dangerous trend.[14] These targeted attacks often slip through signature scanning systems; widespread distribution and publicity are not part of their intended use.

Fraudulent electronic mail attacks are generally referred to as *phishing*, reportedly derived from the combination of phreaking (as in phone phreaking, bypassing long-distance billing systems) and fishing. With some sense of humor, targeted phishing is referred to as *spearphishing*; attacks against senior managers (e.g., CEOs) are referred to as *whaling* or *harpooning*. The nomenclature may be humorous, but well-executed harpooning attacks can be difficult to detect or prevent. Such attacks have been found at the center of several important compromises (e.g., the 2011 spear phishing attack against RSA[15]) at major corporations. High-value targets are not only senior executives; individuals with access to or authority over high-value information are equally at risk.

One of the earliest documented episodes of this type is the 2004 use of Trojan horse software, discussed in §2.10.5 in this *Handbook,* for extensive data theft. This case was not an isolated incident. The trend of targeted attacks against senior personnel, rather than random attacks, has accelerated.[16]

Investigative agencies have also entered the fray. In July 2007, an affidavit filed by the Federal Bureau of Investigation in connection with a series of bomb threats described the use of spyware to infiltrate a suspect's computer and return information to investigators.[17]

Mobile code presents a complex set of issues to information technology (IT) professionals. Allowing mobile code into an operating environment compromises any enclave; however, even commercial off-the-shelf (COTS) programs breach the integrity of an enclave. The differences between mobile code and other forms of code are primarily in the way these external programs are obtained, installed, documented, and controlled. In an enterprise computing environment, COTS acquisition normally involves conscious and explicit evaluation of the costs and benefits of installing a particular named and documented program. In contrast, mobile code programs are installed largely without notification to the user, and generally without documentation, change control, or review of any kind. Unless client-system firewalls are set to reject mobile code automatically (which is increasingly difficult, if not a technical and practical impossibility), system administrators cannot be certain exactly which software has been executed on client machines under their nominal control. The use of Secure Sockets Layer (SSL)–based technologies such as HTTPS to otherwise secure connections used by applications also has the side effect of preventing the detection of mobile code at the firewall level. Although such control is often illusory, due to user circumvention of restrictions on installation of unauthorized software, the use of mobile code, installed by external Web sites, seriously compromises any remaining control over employee software configurations.

Mobile code has also been used in some cases to enforce proprietary rights in content, as was the case in a 2005 affair involving Sony Music.[18] The Sony Music case involved software contained on music CD-ROMs; precisely the same effect could have occurred with a downloaded file or Web page. The covert installation of software of any kind is a serious hazard to integrity, security, and privacy. Malfunction or misuse of such software would likely fit within the criminal statutes defining illegal, unauthorized alteration of systems. The attorney general of Texas,[19] as well as private class actions in New York[20] and California,[21] all filed cases against Sony. All of these actions were settled by Sony BMG in December 2006.[22] In January 2007, the U.S. Federal

Trade Commission announced a settlement with Sony BMG on the charges of installing software without permission.[23] Investigations were also opened by the attorneys general of Massachusetts and Florida, as well as overseas in Italy and Canada.

There are also reports that the Sony Root Kit was exploited by the Backdoor.IRC.Snyd.A exploit[24] and others to hide files from malware scans. The widespread nature of this induced vulnerability should also give pause. A widespread vulnerability provides an ecological niche ready for exploitation.

17.1.2 Motivations and Goals. The motivations and goals of malware propagators have continued an evolutionary trend from the unintentionally destructive to the vengeful, vindictive, and criminal. Unsurprisingly, the actors have migrated from rogue individual pranksters to criminals and nation states and proxies.

In the beginning, many incidents were randomly damaging, or pranks with unintended side effects. This is no longer the case. Now malevolent mobile code is often code with a purpose. That purpose may be embarrassment, it may be blackmail, it may be corporate espionage, or it may be out-and-out theft. In a different dimension, the goal may be the subordination of otherwise innocent computer resources for a criminal enterprise against unrelated third parties.

The change in goals also has a dramatic impact on counter strategies. When the goal was mass publicity, the same infection was widespread, and scanning technologies could be used to identify known threats. When the goal is no longer publicity, publicity and widespread infection are maladaptive. Covert infection is then a far more attractive strategy than mass distribution. Custom mobile code designed to achieve selective covert infections is unlikely to quickly appear in the crosshairs of scanning software. This follows the evolutionary trajectory common in the biological world, where pathogens tend to mutate into less fatal forms over time. It is maladaptive for a parasite, which is what most malware is, to fatally damage its host. The downside of this effect is that the chronic infections with no apparent side effects are often overlooked.

Going forward, technologies and operational routines that make it difficult for unauthorized code to take up residence in or compromise the persistent state of the system are far more desirable counterstrategies than approaches based on scanning for known infections.

17.1.3 Design and Implementation Errors. Design and implementation errors take a variety of forms. The simplest cases involve software that malfunctions on a constant predictable basis. More pernicious and more dangerous are those errors that silently compromise the strict containment of a multiuser environment. Errors in such prophylactic layers, known as brick walls or sandboxes, compromise the integrity of the protection scheme. In the worst cases, they permit unfettered access to system-level resources by unprivileged user programs.

Design and implementation errors can occur within any program or procedure, and mobile code is no exception. *Sandboxes* (nonprivileged, restricted operating environments) are intended to prevent unauthorized operations. *Authentication* determines which organization takes responsibility for such errors.

This chapter looks at a security model based on authentication of mobile code and then examines how restricted operating environments help to limit damage from harmful mobile code.

These concerns are appropriate to both widely distributed and targeted attacks. The challenge of targeted attacks lies in their small population; targeted attacks are unlikely to appear on the radar of general distribution scanning programs.

17.2 SIGNED CODE. Authentication technologies are designed to ensure that information supplied by an organization has not been altered without authorization. The technology used to implement this assurance is based on use of the Public Key Infrastructure (PKI), as discussed in detail in Chapter 37 in this *Handbook*. Code authenticated using PKI-based mechanisms often is referred to as *signed.* Signed code generally is immune to unauthorized modification; however, a signature guarantees integrity only from the point in time that the code is signed; the signing process does not imply safety or quality of the code prior to the point of signing.

Once signed, a file cannot be altered without invalidating the original signature; without the cooperation of someone holding access to the private key associated with the creating organization's X.509 certificate, the signature cannot realistically be tweaked to look legitimate. (An X.509 certificate, digitally signed by an authorized user, authenticates the binding between a user's name and the user's public key.) However, looking below the surface, such precautions do not address a variety of vulnerabilities:

- Unauthorized access to private (signing) keys
- Access to the code base prior to signing
- Fraudulent certificates
- Design and implementation errors

17.2.1 Authenticode. Microsoft's Authenticode technology is an example of an authentication-based approach.[25] Developers wishing to distribute code obtain an appropriate digital certificate from a Certification Authority (CA) and use the digital certificate to sign the code. The signature is then checked by the client system each time that the code is executed.

Authenticode relies on several components:

- PKI and the X.509 certificates issued by a Certification Authority.
- Limited access to the private keys associated with the issuing organization's X.509 certificate. In Microsoft terminology, the term Software Publishing Certificate or SPC refers to a PKCS #7 object, which in turn contains a collection of X.509 certificates used to sign code.
- The integrity of the processes used by the CA to ensure that requests for X.509 certificates are legitimate.

Authenticode does not address issues relating to the safety or accuracy of signed code, merely that it is authentic and unaltered since signing. For example, signing does not provide any guard against employee malfeasance.

17.2.2 Fundamental Limitations of Signed Code. Signing technologies, regardless of the context (e.g., email, applets, and archives), do not directly address questions of accuracy or correctness; they merely address questions of origin. The

biggest danger in signing schemes is the all-or-nothing approach taken to trust. Signed items are presumed to be trustworthy to the fullest extent of the requesting user's authority. The signed item can perform any operation that the user would be permitted to execute. There is no concept of *partial trust*. In an attorney's words, such an acceptance would be a *general power of attorney*. In the words of the CERT/Coordination Center (CERT/CC)–sponsored *Security in ActiveX Workshop*: "A digital signature does not, however, provide any guarantee of benevolence or competence."[26]

At the same time, the inherent power and apparent legitimacy of a digital signature place a heavy burden on signers and the higher levels of the PKI to ensure the integrity of the mechanisms and secrets.

The key to the integrity of signed code is the signing process and the process that generates the object to be signed; the security of the secret keys required for its implementation determines the degree of trust in attribution of the signed code. In the truest sense, the private keys associated with the X.509 certificate represent the keys to the kingdom, as valuable as a signature chop in the Far East or a facsimile signature plate for a bank account.

On a practical level, accepting code signed by an organization is an explicit acceptance that the signing organization has good controls on the use of its signing keys. Organizations that take security seriously, segregating access to privileged accounts and controlling access to systems, are well positioned to manage the procedures for signing code.

Thus, the procedures and systems used for signing code should be treated with the same caution as is used for the aforementioned signing plates or the maximum security cryptographic facilities familiar to those in the national security area.

Unfortunately, despite years of publicity about the dangers of shared passwords and accounts, many IT installations continue to use shared accounts and passwords. There is little reason to assume that the secrets relating to PKI are better protected, despite extensive recommendations that those details be well guarded.

17.2.3 Specific Problems with the ActiveX Security Model. The CERT/CC workshop on Security in ActiveX summarized the security issues in three major areas: importing and installing controls, running controls, and the use of controls by scripts.[27] The key findings from this report are as follows:

- Importing and Installing Controls
  - As discussed, the sole basis for trusting a signed control is its presumed origin. However, the originator of the code may have incorporated a design flaw in the control or may not have done adequate software quality assurance to prevent serious bugs.
  - A trusting user may install a signed control that contains a vulnerability making it useful for attackers simply because it is signed.
  - On Windows systems with multiple users, once a control has been permitted by one user, it remains available for all users, even if their security stances differ.
- Running Controls
  - An ActiveX control has no limitations on what it can do on the client machine, and it runs with the same privileges as those of the user process that initiated the control.

- Although ActiveX security measures are available in Internet Explorer, other client software may run controls without necessarily implementing such security. Internet Explorer security levels tend to be all or nothing, making it difficult to allow a specific control without allowing all controls of that type. Remote activation of controls can bypass normal security perimeters such as those imposed by firewalls.
- There is no basis for deciding whether a particular control is safe to execute or not, in any particular context.

- Scripting Concerns
  - Lacking a general basis for limiting the actions of controls, ActiveX programmers must effectively determine their own precautions to prevent harmful actions. It is difficult enough to develop a good set of boundaries on program activity, even if one uses a general model such as the sandbox described later; it is extremely difficult to see how individual developers can be expected or, more importantly, trusted to create their own equivalent of the sandbox for each individual control. In light of these hazards, the authors of the CERT/CC report stated that "there is a large number of potential failure points."

17.2.4 Case Studies. Several security breaches or demonstrations mediated through ActiveX have occurred since the introduction of this technology in the mid-1990s.

17.2.4.1 Internet Exploder. In 1996, Fred McLain wrote Internet Exploder, an ActiveX control designed to illustrate the broad degree of trust conferred on an ActiveX control by virtue of its having been signed. Exploder, when downloaded for execution by Internet Explorer, will shut down the browser's computer (the equivalent of the *Shut down* | *Shut down* sequence from the Start menu on a Windows system). This operation is operationally disruptive but does not actually corrupt the system. McLain notes in his frequently asked questions (FAQ) on Exploder that it is easy to build destructive or malicious controls.[28]

Exploder raises an important question: Who and what are the limits on trust when using signed code? In normal commercial matters, there is a large difference between an inauthentic signature, a forgery, and a properly signed but unpayable check. In software, the difference between an inauthentic control and a dangerous one is far less clear.

17.2.4.2 Chaos Computer Club Demonstration. On January 27, 1997, a German television program showed members of the Chaos Computer Club demonstrating how they could use an ActiveX control to steal money from a bank account. The control, available on the Web, was written to subvert the popular accounting package Quicken using what amounts to an electronic version of tailgating. A victim need merely visit a site and download the ActiveX control in question; the control then automatically checked to see if Quicken was installed. If so, the control ordered Quicken to issue a transfer order to be saved in its list of pending transfers. The next time the victim connected to the appropriate bank and sent all pending transfer orders to the bank, all the transfers would be executed as a single transaction. The user's personal identification number (PIN) and transaction authorization number (TAN) would apply to all the transfers, including the fraudulent one in the pile of orders. Most

victims would be unaware of the theft until they received their next statement—if then.[29]

Dan Wallach of Princeton University, commenting on this case, wrote:

> When you accept an ActiveX control, you're allowing completely arbitrary code to rummage around your machine and do anything it pleases. That same code could make extremely expensive phone calls, to 900 numbers or over long distances, with your modem; it can read, write, and delete any file on your computer; it can install Trojan horses and viruses. All without any of the subterfuge and hackery required to do it with Java. ActiveX hands away the keys to your computer.[30]

Responding to criticisms of the ActiveX security model, Bob Atkinson, architect and primary implementer of Authenticode, wrote a lengthy essay explaining his point of view. Among the key points:

- Microsoft never claimed that it would certify the safety of other people's code.
- Authentication is designed solely to permit identification of the culprits *after* malicious code is detected.
- Explorer-based distribution of software is no more risky than conventional purchases through software retailers.[31]

Subsequent correspondence in the RISKS Forum chastised Mr. Atkinson for omitting several other key points, such as:

- Interactions among ActiveX controls can violate system security even though individual controls appear harmless.
- There is no precedent in fact for laying liability at the feet of software developers even when you can find them.
- Under attack, evidence of digital signature is likely to evaporate from the system being damaged.
- Latency of execution of harmful payloads will complicate identification of the source of damage.
- Malice is not as important a threat from code as incompetence.
- Microsoft has a history of including security-threatening options, such as automatic execution of macros in Word, without offering any way of turning off the feature.
- A Web site can invoke an ActiveX control that is located on a different site or that already has been downloaded from another site, and can pass, by means of that control, unexpected arguments that could cause harm.[32]

Wallach's criticisms of ActiveX are generally applicable to signed code technologies, not peculiar to Microsoft's ActiveX.

17.2.4.3 Certificates Obtained by Imposters. In January 2001, VeriSign issued two *Class 3 Digital Certificates* for signing ActiveX controls and other code to someone impersonating a Microsoft employee. As a result, users receiving code signed using these certificates would receive a request for acceptance or rejection of a certificate apparently signed by Microsoft on January 30 or 31, 2001. As Russ Cooper

commented on the NTBUGTRAQ Usenet group when the news came out in March 2001:

> The fact that unless you actually check the date on the Certificate you won't know whether or not its [*sic*] one you can trust is a Bad Thing(tm)[sic], as obviously not everyone (read: next to nobody) is going to check every Certificate they get presented with.
>
> You gotta wonder how VeriSign's issuance mechanism could be so poorly designed and/or implemented to let something like this happen.
>
> Meanwhile, Microsoft are [*sic*] working on a patch that will stick its finger in this dam.
>
> Basically, VeriSign Code-Signing Certificates do not employ a Certificate Revocation List (CRL) feature called CDP, or CRL Distribution Point, which causes the Certificate to be checked for revocation each time its read. Even if you have CRL turned on in IE, VeriSign Code-Signing Certificates aren't checked.
>
> Microsoft's update is going to shim in some mechanism which causes some/all Code-Signing Certificates to check some local file/registry key for a CRL, which will (at least initially) contain the details of these Certificates. Assuming this works as advertised, any attempt to trust the mis-issued Certificates should fail.[33]

Roger Thompson, Chief Technical Officer for Exploit Prevention Labs, explained that the imposters' motives would determine how bad the results would be from the fraudulent certificates. "If it was someone with a purpose in mind, then six weeks is a long time to do something," he said. "If the job was to install a sniffer, then there could be a zillion backdoors as a result of it." Published reports indicated that the failure of authentication occurred due to a flaw in the issuing process at VeriSign: The certificates were issued *before* receiving verification by email that the official customer contact authorized the certificates. This case was the first detected failure of authentication in over 500,000 certificates issued by VeriSign.[34]

Some recent cases involving stolen or forged certificates include the following:

- In November 2011, antimalware vendor F-Secure's Mikko Hypponen reported "finding a piece of malicious software that was cryptographically signed by a forged Adobe certificate originating with Government of Malaysia: Malaysian Agricultural Research and Development Institute …"[35]
- In September 2012, Adobe reported that they "… received two malicious utilities that appeared to be digitally signed using a valid Adobe code signing certificate."[36] The company explained, "Sophisticated threat actors use malicious utilities like the signed samples during highly targeted attacks for privilege escalation and lateral movement within an environment following an initial machine compromise."
- In February 2013, Bit9 Systems revealed that "Due to an operational oversight within Bit9, we failed to install our own product on a handful of computers within our network. As a result, a malicious third party was able to illegally gain temporary access to one of our digital code-signing certificates that they then used to illegitimately sign malware. There is no indication that this was the result of an issue with our product. Our investigation also shows that our product was not compromised."[37]

17.2.4.4 Certificate Forgery. Certificate forgery is another hazard. Certificates are not magical; they are merely documented cryptographically signed by a recognized Certification Authority. The cryptographic techniques used to sign certificates

are subject to attack. A research result demonstrating the hazard of MD-5 collisions leading to forged certificates was originally published in 2008.[38] For this reason, many authorities have recommended that MD-5 signed certificates be withdrawn from use. In August 2011, an MD-5 signed fraudulent certificate for ***.google.com** was reported by Iranian Internet users.[39]

17.3 RESTRICTED OPERATING ENVIRONMENTS. From a Web perspective, the term *sandbox* defines what could be referred to as a restricted operating environment. Restricted operating environments are not new; they have existed for nearly 50 years in the form of multiuser operating systems, including MULTICS, OS/360 and its descendants, OpenVMS, UNIX, and others.

In simple terms, a restricted, or nonprivileged, operating environment prohibits normal users and their programs from executing operations that can compromise the overall system. In such an environment, normal users are prohibited from executing operations such as HALT that directly affect hardware. User programs are prevented from executing instructions that can compromise the operating system memory allocation and processor state and from accessing or modifying files belonging to the operating system or to other users. Implemented and managed carefully, such systems are highly effective at protecting information and data from unauthorized modification and access. The National Computer Security Center (NCSC) *Orange Book* contains criteria for classifying and evaluating trusted systems.[40]

The strengths and weaknesses of protected systems are well understood. Permitting ordinary users unrestricted access to system files compromises the integrity of the system. Privileged users (i.e., those with legitimate access to system files and physical hardware) must be careful that the programs they run do not compromise the operating system. Most protected systems contain a collection of freestanding programs that implement useful system functions requiring some form of privilege to operate. Often these programs have been the source of security vulnerabilities. This is the underlying reasoning behind the universal recommendation that programs not run as root or Administrator or with similar privileges unless absolutely necessary.

17.3.1 Java. Java is a language developed by Sun Microsystems for platform independent execution of code, typically within the context of a Web browser. The basic Java environment includes a Java Virtual Machine (JVM) and a set of supporting software referred to as the Java Run Time Environment. Applets downloaded via the World Wide Web (intranet or Internet) have strict limitations on their ability to access system resources. In particular, these restrictions prevent the execution of external commands and read or write access to files.

The Java environment does provide for *signed applets* that are permitted wider access to files. Dynamically downloaded applets also are restricted to initiating connections to the system that supplied them, theoretically limiting some types of third-party attacks.

In concept, the Java approach, which also includes other validity tests on the JVM pseudocode, should be adequate to ensure security. However, the collection of trusted applets found locally on the client system and signed downloaded applets represent ways in which the security system can be subverted. Without signature, the Java approach is also vulnerable to attack by Domain Name System (DNS) spoofing.

Multiuser protection and virtual machine protection schemes also are totally dependent on the integrity of the code that separates the nonprivileged users from privileged, system-compromising, operations. Java has not been an exception to this rule. In 1996,

a bug in the Java environment contained in Netscape Navigator Version 2.0 permitted connections to arbitrary Universal Resource Locators (URLs).[41] Later, in 2000, errors were discovered in the code that protected various resources.[42] Although the Java environment is less widely exploitable than ActiveX, vulnerabilities continue to be uncovered. In 2007, at least two vulnerabilities were reported by US-CERT.[43] Significantly, both of these reported vulnerabilities involved the ability of untrusted applets to compromise the security envelope.

Additionally, since unsigned code can take advantage of errors in underlying signed code, there is no guarantee that complex combinations of untrusted and trusted code will not lead to security compromises.

17.3.2 Virtual Machines. Isolation can also be realized using virtual machines. Theoretically, each virtual machine runs its own copy of an operating system on a separate virtual copy of the hardware. Properly implemented, mainframe systems have run multiple, independent machines on a single physical system for decades. Virtual machine technology is a foundational technology for so-called cloud computing. In such configurations, the Achilles' heel is the layer that separates the different virtual machines. If there is a deficiency in that layer, or in its interfaces, the compartmentalization may be subject to compromise.

In 2012, just such an incident occurred with multiple virtual machine implementations implementing virtual machines on Intel's x86 processors.[44] In these cases, a specially crafted stack frame could lead to a break of the virtual machine's containment mechanisms. Virtual machines from Xen, FreeBSD, Red Hat, and Microsoft Windows were subject to this attack.

17.3.3 Operating Environment Summary. The details of virtualization are significant in this context. Errors in virtualization software implementation can create security compromises in the separation between different virtual machines, or indeed it may compromise the underlying integrity of the host. Such hazards are not theoretical; they occur in application-level sandboxes and in systems-level virtual machines.

17.4 DISCUSSION. Mobile code security raises important issues about how to handle relationships in an increasingly interconnected computing environment.

17.4.1 Asymmetric, and Transitive or Derivative, Trust. It is common for cyberrelationships to be *asymmetric* with regard to the size or power of the parties. This fact increases the potential for catastrophic interactions. It also creates opportunities for mass infection across organization boundaries. Large or critical organizations often can unilaterally impose limitations on the ability of partner organizations to defend their information infrastructure against damage.

The combination of a powerful organization and insufficient controls on signing authority, or, alternatively, the obligatory execution of unsigned (or self-signed) ActiveX controls, is a recipe for serious problems. The powerful organization is able to obligate its partners to accept a low security level, such as would result, for example, from using unsigned ActiveX controls, while abdicating responsibility for the resulting repercussions.

All organizations should, for security and performance reasons, use the technology that requires the least degree of privilege to accomplish the desired result. JavaScript/ECMAScript can provide many functions, without the need for the

functionality provided by Java, much less ActiveX. It remains common for large organizations to force the download of ActiveX controls for purposes that do not require the power of ActiveX, merely using the justification that they perceive Internet Explorer to be the more prevalent browser. Often this requires running the installation script from an account with Administrator privileges, a second security violation. This is particularly surprising since these same organizations often offer parallel support for Firefox, Opera, Safari, and other non-ActiveX supporting browsers on Linux, Apple, and other platforms. This Trust Me concept forces the risk and burden of consequences on the end user, who is far less able to deal with the consequences.

As noted earlier in this chapter, Web servers represent an attractive vector for attacks. Signing (authentication) methods are a way to control damage potential, if the mechanisms used for admitting executable code are properly controlled. Failure to control these mechanisms leads to severe side effects.

In Chapter ***30*** in this *Handbook*, it is noted that protecting Web servers requires that the server contents be managed with care. It is appropriate and often necessary to isolate Web servers on separate network segments, separated from both the Internet and the organizational intranet by firewalls. These precautions are even more necessary when servers are responsible for supplying executable code to clients.

Security practitioners should carefully examine the different functions performed by each server. In some cases, such as OpenVMS hosts, where network servers commonly run as unprivileged processes in separate contexts and directory trees, it is feasible to run multiple services on a single server. In other systems, such as UNIX and Windows, where it is common for applications services to execute as privileged, with full access to all system files, a logic error in a network service can compromise the security of the entire server, including the collection of downloadable applets.

Far more serious and equally subtle is *transitive* (or derivative) trust: Alpha trusts Beta who trusts Gamma. A security compromise—for example, an unsigned Java applet or a malfunctioning or malevolent ActiveX control supplied by Gamma—compromises Beta. Beta then causes problems with Alpha's systems. This cascade can continue repeatedly, leading to numerous compromised Web services and systems far removed geographically and organizationally from the original incident.

17.4.2 Misappropriation and Subversion. The threat space has mutated over the last several years. Where the main danger from mobile code was attacks on the target machine, today's threat is far more diverse. In November 2007, John Schiefer of Los Angeles pled guilty to installing software designed to capture usernames and passwords. According to news reports, he was also involved in running a number of networks of compromised computers, often referred to as bots, which are often used to initiate distributed denial-of-service (DDoS) and other attacks.[45] In this particular case, the announcement by the U.S. Department of Justice[46] mentions two specific episodes: 250,000 machines infected with spybots to obtain user's usernames and passwords for PayPal and other systems; and a separate scheme involving a Dutch Internet advertising company in which a network of 150,000 infected computers were used to sign up for one of the advertising company's programs.

This was one of the cases stemming from Bot Roast II,[47] an FBI operation against several botnet networks.

17.4.3 Multidimensional Threat. Mobile code is a multidimensional threat, with several different aspects that must each be treated separately. Signing code, such as Java applets or ActiveX controls, addresses the problem of authenticity and authority

to release the code. However, the integrity of the signature mechanism requires that the integrity of the PKI infrastructure be beyond reproach. In a very real sense, the PKI infrastructure is beyond the control of the organization itself. Any compromise or procedural slip on the part of the Certificate Authority or signer invalidates the presumptions of safety.

Signing, however much it contributes to resolving the question of authenticity, does not address safety or validity. As an example, the Windows Update ActiveX control, distributed by Microsoft as part of the various Windows operating systems, has as its underlying purpose the update of the operating system. A failure of that control would be catastrophic. Fortunately, Microsoft gives users the choice of using the automatic update facility or doing updates manually. Many Web applications are not so accommodating.

The problem is not solely a question of malfunctioning applets. It is possible that a collection of applets involved in a client's overall business activities may collide in some unanticipated fashion, from attempting to use the same Windows registry key in contradictory ways, to inadvertently using the same temporary file name. Similar problems often occur with applications that presume they have a monopoly on the use of the system, an all-too-common syndrome.

These issues are, for the most part, completely unrelated to each other. A solution in one area would neither improve nor worsen the situation with regard to the other issues.

17.4.4 Client Responsibilities. The expanding threat presents a challenge to those responsible for ensuring the integrity of desktop computing. Put simply, there is a complex, multidimensional threat, and it is not easily defended against using the techniques of portals, firewalls, and scanners.

The danger from browsing the World Wide Web is the danger that the browser will permit an attacker, directly or indirectly, to cause a modification to the persistent state of the system. The simplest step in the correct direction is not to browse the World Wide Web from within a protection context that has access to critical system files and settings. Limiting this access by using a nonprivileged user account for browsing significantly decreases the hazard, provided of course that the system files are protected from access by such an account.

The mass availability of *virtual machine* technology presents an additional alternative. Virtual machine technology, pioneered by IBM in the 1960s on mainframes, has emerged in a new guise on platforms down to the desktop level. The general availability of this capability in the desktop world opens up a whole new defensive strategy against mobile malware: the expendable Web browser.

An expendable WWW browser is an instantiated desktop within a virtual machine environment, from a known system image. If it is compromised, it is merely rewritten from a known, uncompromised system image. It allows one to create a low-security, at-risk, browsing enclave within an otherwise higher security environment. This is an approach that has been used, in a physical sense to be sure, by some organizations proving public access personal computers. Rather than attempting to fortify the machines against compromise or attach, they are reinitialized from a known image after each user. This allows the end user to indulge the foibles of trading partners' attempts to impose unsafe computing practices in an expendable environment that can be isolated. Using Windows as an example, while it is an unsafe practice to install software as Administrator, it is far less damaging to do so in a virtual machine, where the machine can be deleted at convenience with little side effect.

Similarly, *disposable virtual machines*[48,49] can be used by normal users to separate the normal operating environment from one where there is a perceived increase in risk exposure (e.g., browsing suspect sites; forced installation of ActiveX controls of uncertain provenance or safety).

17.4.5 Server Responsibilities. As noted earlier in this chapter, Web servers represent an attractive vector for attacks. Signing (authentication) methods are a way to control damage potential, provided the mechanisms used for admitting executable code are properly controlled. Failure to control these mechanisms leads to severe side effects.

The concept of *minimum necessary privilege* applies to mobile code. There is little reason to impose the use of ActiveX for the purposes of changing the color of a banner advertisement. This effect can often be accomplished using Cascading Style Sheets (CSS). JavaScript/ECMAScript is capable of many powerful, display-related operations with a high degree of safety in situations where CSS is not adequate. Using Java to maintain a shopping cart (price, quantities, and contents) is reasonable and does not require the use of a signed applet, with its attendant greater capabilities and risks. At the other end of the scale, it is plausible that a system update function (e.g., the Windows Update function, which automatically downloads and installs changes to the Windows operating system) requires the unbridled power of a signed ActiveX control.

When the power of signed applets or controls is required, good software engineering practice provides excellent examples of how to limit the potential for damage and mischief, as discussed in Chapter 38 in this *Handbook*.

These problems extend beyond an individual organization's direct exposure. An increasing trend over the past several years has been to use third-party Web servers as an infection vector, a springboard for infecting Website visitors. If components are downloaded from third-party servers, the integrity of your WWW site may be vulnerable to compromise.

Good software implementation isolates functions and limits the scope of operations that require privileged access or operations. Payroll applications do not directly manipulate printer ports, video display cards, network adapters, or disk drives. Privileged operating system components, such as device drivers and file systems, are responsible for the actual operation. This separation, together with careful parameter checking by the operating system kernel and the privileged components, ensures safety.

The same techniques can be used with applets and controls. Because they require more access, they should be programmed carefully, using the same defensive measures as are used when implementing privileged additions to operating systems. As an example, there is little reason for a Simple Mail Transfer Protocol (SMTP) server to be privileged. An SMTP server requires privileges for a single function, the delivery of an individual email message to a recipient's mailbox. This can be accomplished in two ways:

1. Implement the application in a nonprivileged way, by marking users' email files and directories with the necessary access permissions for the mail delivery program to create and modify email files and directories. Such a mechanism is fully in conformance with the NCSC *Orange Book*'s C2-level of security.
2. Implement a separate subcomponent whose sole responsibility is the actual message delivery of the message. The subcomponent must be written defensively to check all of its parameters, and does not provide an interface for the execution of arbitrary code. This approach is used by HP's OpenVMS operating system.

The UNIX *sendmail* program, by contrast, is a large, multifunctional program that executes with privileges. *sendmail* has been the subject of numerous security problems for over a decade and has spawned efforts to produce more secure replacements.[50]

17.5 SUMMARY. Mobile code provides many flexible and useful capabilities. The different mechanisms for implementing mobile code range from the innocuous (HTML), to fairly safe (JavaScript/ECMAScript), and with increasing degrees of power and risk through Java and ActiveX.

Ensuring security and integrity with the use of mobile code requires cooperation on the part of both the provider and the client. Clients should not accept random signed code and controls. Providers have a positive responsibility to:

- Follow good software engineering practices
- Grant minimum necessary privileges and access
- Use defensive programming
- Limit privileged access, with no open-ended interfaces
- Ensure the integrity of the signing process and the associated private keys

With appropriate caution, mobile code can be a constructive, powerful part of intranet and Internet applications, both within an organization and in cooperation with its customers and other stakeholders.

17.6 FURTHER READING

Carl, Jeremy. "ActiveX Security: Under the Microscope," *Web Week*, 2, No. 17, November 4, 1996; www.Webdeveloper.com/activex/activex_security.html

CERT. "NIMDA Worm," September 11, 2001, www.cert.org/advisories/CA-2001-26.html

CERT. "sadmind/IIS Worm," May 8, 2001, www.cert.org/advisories/CA-2001-11.html

CERT. "Unauthentic 'Microsoft Corporation' Certificates," March 22, 2001, www.cert.org/advisories/CA-2001-04.html

Dormann, Will, and Jason Rafail. "Securing Your Web Browser," CERT, January 23, 2006, www.cert.org/tech_tips/securing_browser/index.html on March 19, 2013.

Evers, J. FAQ: JavaScript Insecurities. ClNET July 28, 2006 < http://news.cnet.com/2100-7349_3-6100019.html >

Felten, Edward. "Security Tradeoffs: Java vs. ActiveX," last modified April 28, 1997, www.cs.princeton.edu/sip/faq/java-vs-activex.html

Felten, E., and J. Halderman (2006, February 14). *Lessons from the Sony CD DRM Episode*, Center for Information Technology Policy, Department of Computer Science, Princeton University, USENIX Security 2006 http://static.usenix.org/events/sec06/tech/full_papers/halderman/halderman.pdf

Felten, E., and G. McGraw. *Securing Java: Getting Down to Business with Mobile Code*. New York: John Wiley & Sons, 1999. Also free and unlimited Web access from www.securingjava.com

Gehtland, J., B. Galbraith, and D. Almaer. *Pragmatic Ajax*. Raleigh, NC: Pragmatic Bookshelf, 2006.

Grossman, J., and T. C. Niedzialkowski. "Hacking Intranet Websites from the Outside," *Black Hat (USA)*, Las Vegas, August 3, 2006. www.blackhat.com/presentations/bh-usa-06/BH-US-06-Grossman.pdf

Hensing, R. "W32/HLLP.Philis.bq, Chinese Gold Farmers and What You Can Do about It," December 2, 2007, http://blogs.technet.com/robert_hensing/archive/2006/12/04/w32-hllp-philis-bq-chinese-gold-farmers-and-what-you-can-do-about-it.aspx on

Holzman, S. *Ajax Bible*. Hoboken, NJ: John Wiley & Sons, 2007.

Keizer, G. "FBI Planted Spyware on Teen's PC to Trace Bomb Threats," *PCWorld*, July 19, 2007. www.pcworld.com/article/134823/article.html

McGraw, G., and E. W. Felten. *Securing Java: Getting Down to Business with Mobile Code*, 2nd Edition. John Wiley & Sons, 1999.

McGraw, Gary & Ed Felten. "Java Security Web Site." www.cigital.com/javasecurity/index.html

Microsoft. "Introduction to Code Signing" (with appendix), 2001, http://msdn.microsoft.com/en-us/library/ms537361%28v=vs.85%29.aspx

Rhoads, C. "Web Scammer Targets Senior U.S. Executives," *Wall Street Journal*, November 9, 2007.

Schwartz, J. "iPhone Flaw Lets Hackers Take Over Security Firm Says," *New York Times*, July 23, 2007.

VeriSign, "Microsoft Security Bulletin MS01-017: Erroneous VeriSign-Issued Digital Certificates Pose Spoofing Hazard," March 22, 2001, www.microsoft.com/TechNet/security/bulletin/MS01-017.asp

VeriSign, "VeriSign Security Alert Fraud Detected in Authenticode Code Signing Certificates," March 22, 2001.

Zakas, N., J. McPeak, and J. Fawcett. *Practical Ajax*, 2nd ed. Hoboken, NJ: John Wiley & Sons, 2007.

17.7 NOTES

1. ECMA-262 retrieved from www.ecma-international.org/publications/files/ECMA-ST-ARCH/ECMA-262%205th%20edition%20December%202009.pdf
2. Carter Dotson, "Is Apple Losing Its Grip on App Store Security?" Mashable Web site, March 1, 2012, http://mashable.com/2012/03/01/app-store-security-risks
3. Scott Thurm and Yukari Iwatani Kane, "Your Apps Are Watching You," *The Wall Street Journal*, December 17, 2010, http://online.wsj.com/article/SB10001424052748704694004576020083703574602.html
4. J. Schwartz, "iPhone Flaw Lets Hackers Take over, Security Firm says," *The New York Times*, July 23, 2007, www.nytimes.com/2007/07/23/technology/23iphone.html
5. CBC News (2013, January 28) "WhatsApp breaches privacy laws," *CBC News*, January 28 2013, www.cbc.ca/news/technology/story/2013/01/28/technology-whatsapp-privacy.html
6. Office of Privacy Commissioner of Canada, "Report of Findings Investigation into the personal information handling practices of WhatsApp Inc.," January 15, 2013, www.priv.gc.ca/cf-dc/2013/2013_001_0115_e.asp
7. Memorandum by Arthur M. Money, Assistant Secretary of Defense for C3I, and CIO to Secretaries of the Military Departments, Chairman of the Joint Chiefs of Staff, Chief Information Officers of the Defense Agencies, et al., "Use of Mobile Code Technologies in DoD Information Systems," October 23, 2006, www.dtic.mil/whs/directives/corres/pdf/855201p.pdf

8. J. Penny, "40 Years Too Long in Norwich Porn Case?" *Norwich Bulletin*, January 9, 2007; G. Smith, "Teacher Facing Porn Charges," *Norwich Bulletin*, November 11, 2004.
9. Robert Cringley "The Julie Amero Case: A Dangerous Farce," *PCWorld*, December 2, 2008, http://pcworld.com/article/154768/julie_amero.html
10. Robert Cringley, "Julie Amero Case, Part II: May the Farce Be With You," *PCWorld*, December 4, 2008, www.pcworld.com/article/154937/julie_amero.html
11. Robert McMillan, "How Spyware Nearly Sent a Teacher to Prison," *PCWorld*, November 26, 2008, www.pcworld.com/article/154611/spyware_teacher_prison.html
12. Jesse James Garrett, "AJAX: A New Approach to Web Applications," *Adaptive Path*, February 18, 2005, http://adaptivepath.com/ideas/essays/archives/000385.php
13. Ian Fette and Alexey Melnikov, "The WebSocket Protocol," RFC 6455, December 2011, www.rfc-editor.org/rfc/rfc6455.txt
14. D. Izenberg, "Trojan Horse Developers Indicted," *Jerusalem Post*, March 5, 2006.
15. Sean Gallagher, "RSA Details March cyberattack, Blames 'Nation State' for SecurID Breach," *ars technica*, October 12, 2011, http://arstechnica.com/business/2011/10/rsa-details-march-cyber-attack-blames-nation-state-for-securid-breach
16. J. Kirk, "Hackers Target C-level Execs and Their Families," *Network World*, July 2, 2007, www.networkworld.com/news/2007/070207-hackers-target-c-level-execs-and.html
17. G. Keizer, "FAQ: What We Know (Now) about the FBI's CIPAV Spyware," *Computerworld*, July 29, 2007, www.computerworld.com/s/article/9028298/FAQ_What_we_know_now_about_the_FBI_s_CIPAV_spyware
18. Mark Russinovich, "Sony, Rootkits and Digital Rights Management Gone Too Far," Mark Russinovich's Blog, December 2, 2007, http://blogs.technet.com/markrussinovich/archive/2005/10/31/sony-rootkits-and-digital-rights-management-gone-too-far.aspx
19. *The State of Texas v. Sony BMG Music Entertainment, LLC*, Case GV-505065, District Court of Travis County, Texas, 126th Judicial District.
20. *James Michaelson and Ori Edelstein v. Sony BMG Music, Inc.* and First 4 Internet, Case 05 CV 9575, United States District Court, Southern District of New York.
21. *Alexander William Guevara v. Sony Music Entertainment, et al.*, Case BC342359, Superior Court of the State of California, County of Los Angeles.
22. R. McMillan, "Sony Pays $ 1.5M to Settle Texas, California Root Kit Suits," *Computerworld*, December 20, 2006.
23. Federal Trade Commission, "Sony BMG Settles FTC Charges," January 30, 2007, www.ftc.gov/opa/2007/01/sony.htm
24. Sorin Dudea, "Backdoor.IRC.Snyd.A," Bitdefender Toolbox Web site, December 2, 2007, www.bitdefender.com/VIRUS-1000058-en–Backdoor-IRC-Snyd-A.html
25. Roger Grimes, "Authenticode," Microsoft TechNet Web site, n.d., http://technet.microsoft.com/en-us/library/cc750035.aspx
26. CERT/CC, *Results of the Security in ActiveX Workshop, Pittsburgh, Pennsylvania, August 22–23, 2000*, www.cert.org/archive/pdf/activeX_report.pdf
27. CERT/CC, pp. 6–9.

28. F. McLain, "The Exploder Control Frequently Asked Questions (FAQ)," January 30, 1997, updated February 7, 1997, www.halcyon.com/mclain/ActiveX/Exploder/FAQ.htm
29. D. Weber-Wulff, "Electronic Funds Transfer without Stealing PIN/TAN," *RISKS* 18, No. 80 (1997), http://catless.ncl.ac.uk/Risks/18.80.html
30. D. Wallach, "RE: Electronic Funds Transfer without Stealing PIN/TAN," *RISKS* 18, No. 81 (1997), http://catless.ncl.ac.uk/Risks/18.81.html
31. B. Atkinson, "Comments and Corrections Regarding Authentication," *RISKS* 18, No. 85 (1997), http://catless.ncl.ac.uk/Risks/18.85.html
32. *RISKS* 18, No. 86 (1997), et seq. http://catless.ncl.ac.uk/Risks/18.86.html
33. Microsoft, "Erroneous VeriSign-Issued Digital Certificates Pose Spoofing Hazard," Microsoft Security Bulletin MS01-017, http://support.microsoft.com/kb/293818
34. R. Lemos, "Microsoft Warns of Hijacked Certificates," *ZDNet News*, March 22, 2001, http://news.cnet.com/Microsoft-warns-of-hijacked-certificates/2100-1001_3-254586.html
35. Cory Doctorow, "Stolen Government of Malaysia Certificate Used to Sign Malicious Fake Adobe Software Update," BoingBoing.net Web site, November 14, 2011, http://boingboing.net/2011/11/14/stolen-government-of-malaysia.html
36. Brad Arkin, "Inappropriate Use of Adobe Code Signing Certificate," Adobe Secure Software Engineering Team (ASSET) Blog, September 27, 2012, http://blogs.adobe.com/asset/2012/09/inappropriate-use-of-adobe-code-signing-certificate.html
37. Patrick Morley, "Bit9 and Our Customers' Security," Bit9 Website, February 8, 2013, https://blog.bit9.com/2013/02/08/bit9-and-our-customers-security/
38. Alexander Sotirov, "Creating a Rogue CA Certificate" *Security Research* Web site, December 30, 2008, www.phreedom.org/research/rogue-ca
39. Dan Goodin, "Fraudulent Google Credential Found in the Wild: Did Counterfeit SSL Cert Target Iranians?" *The Register*, August 29, 2011, www.theregister.co.uk/2011/08/29/fraudulent_google_ssl_certificate/
40. For full text, see http://csrc.nist.gov/publications/secpubs/rainbow/std001.txt
41. CERT, "Java Implementations Can Allow Connection to an Arbitrary Host," CERT Advisory, March 5, 1996, updated September 24, 1997, www.cert.org/advisories/CA-1996-05.html
42. CERT, "Netscape Allows Java Applets to Read Protected Resources," CERT Advisory, August 10, 2000, www.cert.org/advisories/CA-2000-15.html
43. CERT, "Sun Java JRE Vulnerable to Unauthorized Network Access," Vulnerability Note VU#336105, October 5, 2007, revised October 12, 2007, www.kb.cert.org/vuls/id/336105; and "Sun Java JRE Vulnerable to Privelege Escalation," Vulnerability Note VU#102289, January 9, 2007, revised May 16, 2007, www.kb.cert.org/vuls/id/102289
44. CERT, "SYSRET 64-Bit Operating System Privilege Escalation Vulnerability on Intel CPU Hardware," Vulnerability Note VU#649219, June 12, 2012, revised September 4, 2012, www.kb.cert.org/vuls/id/649219
45. J. Serjeant, "'Botmaster' Admits Infecting 250,000 Computers," *Reuters*, November 9, 2007.
46. United States Attorney's Office, Central District of California (2007, November 9) "Computer Security Consultant Charges with Infecting up to a Quarter

Million Computers that Were Used to Wiretap, Engage in Identity Theft, Defraud Banks," Press Release No. 07-143, November 9, 2007, www.justice.gov/criminal/cybercrime/press-releases/2007/schieferCharge.pdf

47. FBI, "'Bot Roast II': Cracking Down on CyberCrime," November 29, 2007, www.fbi.gov/news/stories/2007/november/botnet_112907; more extensive details in "'Bot Roast II" Nets 8 Individuals," November 29. 2007. www.fbi.gov/news/pressrel/press-releases/bot-roast-ii-nets-8-individuals
48. Robert Gezelter, "Disposable Virtual Machines: Deliberately Expendable" Ruminations—An IT Blog, August 23, 2010, www.rlgsc.com/blog/ruminations/disposable-virtual-machines.html
49. Robert Gezelter, "Disposable Virtual Machines," 2010 Trenton Computer Festival, Ewing, New Jersey, April 24, 2010, www.rlgsc.com/trentoncomputerfestival/2010/disposable-virtual-machines.html
50. For example, the National Vulnerability Database (NVD, http://nvd.nist.gov) shows that the Common Vulnerabilities and Exposures (CVE) database includes a total of 29 unique vulnerabilities involving sendmail, of which 15 are dated 2000 and 2001. This trend continues, with the NVD Version 2.0 showing an additional 16 sendmail-related issues from 2002 through 2007.

CHAPTER 18

DENIAL-OF-SERVICE ATTACKS

Gary C. Kessler

18.1 INTRODUCTION 18·2

18.2 HISTORY OF DENIAL-OF-SERVICE ATTACKS 18·2

18.3 COSTS OF DENIAL-OF-SERVICE ATTACKS 18·4

18.4 TYPES OF DENIAL-OF-SERVICE ATTACKS 18·6

18.5 SPECIFIC DENIAL-OF-SERVICE ATTACKS 18·6

18.5.1 Destructive Devices 18·6
18.5.2 Email (and Email Subscription) Bombing 18·6
18.5.3 Buffer Overflow 18·8
18.5.4 Bandwidth Consumption 18·8
18.5.5 Routing and Domain Name System Attacks 18·10
18.5.6 SYN Flooding 18·11
18.5.7 Resource Starvation 18·12
18.5.8 Java, PHP, and ASP Attacks 18·13
18.5.9 Router Attacks 18·13
18.5.10 Other Denial-of-Service Attacks 18·13

18.6 PREVENTING AND RESPONDING TO DENIAL-OF-SERVICE ATTACKS 18·14

18.7 DISTRIBUTED DENIAL-OF-SERVICE ATTACKS 18·14

18.7.1 Short History of Distributed Denial of Service 18·15
18.7.2 Distributed Denial-of-Service Terminology and Overview 18·16
18.7.3 Distributed Denial-of-Service Tool Descriptions 18·19

18.8 DENIAL-OF-SERVICE USING EXPLOITABLE SOFTWARE 18·23

18.8.1 Code Red 18·23
18.8.2 NIMDA 18·24

18.9 DEFENSES AGAINST DISTRIBUTED DENIAL-OF-SERVICE ATTACKS 18·25

18.9.1 User and System Administrator Actions 18·25
18.9.2 Local Network Actions 18·26
18.9.3 Internet Service Provider Actions 18·27
18.9.4 Exploited Software Defensive Actions 18·28
18.9.5 Other Tools under Development or Consideration 18·28

18.10 MANAGEMENT ISSUES 18·30

18.11 FURTHER READING 18·31

18.12 NOTES 18·32

18.1 INTRODUCTION. This chapter discusses denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks. These attacks seek to render target systems and networks unusable or inaccessible by saturating resources or causing catastrophic errors that halt processes or entire systems. Furthermore, they are increasingly easy for even *script kiddies* (persons who follow explicit attack instructions or execute attack programs) to launch. Successful prevention of and defense against these attacks will come only when there is widespread cooperation among all Internet service providers (ISPs) and other Internet-connected systems worldwide and when antimalware software is universally installed and kept up to date on user systems.

Working in a variety of ways, the DoS attacker selects an intended target system and launches a concentrated attack against it. Although initially deemed to be primarily a nuisance, DoS attacks can incapacitate an entire network, especially those with hosts that rely on Transmission Control Protocol/Internet Protocol (TCP/IP). DoS attacks on corporate networks and ISPs have resulted in significant damage to productivity and revenues. DoS attacks can be launched against any hardware or operating system platform because they generally aim at the heart of Internet Protocol (IP) implementations. Because IP is the typical target, the DoS attack tools that run under one operating system (Linux is a common choice) can be aimed at any operating system running IP. Additionally, because IP implementations are similar for different platforms, one DoS attack may target several operating systems and work on each. Once written for one platform and released, new DoS attacks appear to evolve (via the examination and participation of hackers and crackers) so that in a short period of time (approximately two weeks) mutations of the DoS attack appear that work on virtually all platforms.

Because of the critical impact that DoS attacks can have, they cannot be taken lightly. Targeted DoS attacks have been around in one form or another since the 1980s; in 1999, they evolved into DDoS attacks, primarily due to the heavy use of internal networks and the Internet. DDoS tools launch coordinated DoS attacks from many sources against one or more targets simultaneously.

This chapter describes first DoS and then DDoS attacks because DoS attacks historically predate DDoS attacks. In addition, some DDoS attacks make use of DoS techniques. However, the attacks themselves, the terminology, and the defenses are sufficiently different that they warrant separate discussion.

18.2 HISTORY OF DENIAL-OF-SERVICE ATTACKS. Historically, any act that prevented use of a system could be called a denial of service. For example, in some mainframe and minicomputer systems, typing a malformed command could cause a system failure; for example, on the HP3000 minicomputer around 1982, using a faulty TERMTYPE parameter in the logon string could instantly crash the system. On the HP3000, typing one or more characters on the system console without pressing the RETURN key would block all further system messages on the console, causing system buffers to fill up with undisplayed messages; when the system buffers (16 by default) were full, no system action requiring notification to the console could take place (e.g., logons, logoffs, special-form requests, or tape-mount requests). A classic DoS without software could be the result of having a cleaner unplug a power cord from a wall socket without realizing that the cord was powering one or more workstations. Configuring a permanent lockout of users after a certain number of bad-password attempts (instead of a timed lockout) led to DoS; malicious users could even try to lockout root users by deliberately entering bad passwords on their accounts. However, such events were usually the result of bugs in the operating system, inadequate numbers of system resources, poor system management, or accident.

One of the first major DoS incidents that made headlines and ripples was probably accidental. It took place in December 1987, when an employee of IBM in Europe sent out a holiday greeting email message. The email message, however, contained a program that drew a nice Christmas tree on the recipient's terminal—and read the recipient's NAMES and NETLOG files listing the user's address book as well as the email addresses of individuals that this user had recently sent mail to or received mail from. The Christmas Tree program then triggered the automatic transmission of copies of itself to all email addresses found in NAMES and NETLOG—and failed to prevent remailing to systems that had already received or sent the greeting. The resulting mailstorm overloaded the IBM corporate network worldwide and brought it crashing down both in Europe and the United States. In addition, messages escaped from IBM's corporate network and wreaked havoc on the BITNET/EARN education and research networks in North America and Europe. Although the cause of the outage was originally thought to be a computer virus, and the incident has been referred to as a virus, a worm, and a prank, the result was a true denial of service.

Perhaps the most famous Internet DoS resulted from the Morris Worm, also referred to as the Internet Worm, which was released in November 1988. Cornell graduate student Robert T. Morris had cowritten an article about vulnerabilities in *Sendmail* and *fingerd* on UNIX systems during the previous year. Most of the TCP/IP community dismissed the vulnerabilities as theoretical; apparently Morris wanted to demonstrate that the vulnerabilities actually could be practically exploited. To demonstrate the problem, his program had to invade another system, guess some passwords, and then replicate itself. Morris said after the incident that he wanted the program to replicate just a few times to demonstrate that it was real, and he included code to prevent rapid spread; unfortunately, a programming error led the worm to replicate often and quickly, in addition to superinfecting already-infected systems. The worm clogged the Internet with hundreds of thousands of messages and effectively brought the entire network down; most sites that did not actually crash were disconnected from the network by their system administrators to avoid being infected and to allow disinfection. Regardless of intent, the Morris Worm inadvertently caused a DoS. Because sites administrators did not have a way to communicate with other admins, and because solutions that Morris himself posted could not be trusted, there was a lot of time spent in eradicating the worm. One result was the creation of Computer Emergency Response Team Coordination Center, now CERT/CC, at Carnegie Mellon University.

On Friday, September 6, 1996, PANIX, a public access Internet provider in Manhattan, was struck by a DoS attack that consisted of many messages flooding the server with massive amounts of data. The attacks were made against mail, news, name, and Web servers as well as user shell account machines. The attackers successfully eluded tracking and the attacks went on for several days. About a week after the attacks began, PANIX went public with the story, and dozens of other online service providers acknowledged that they too were the victims of similar attacks.

In 1997, a disgruntled former employee of Forbes, Inc., used a former colleague's password to remotely access the firm's computer systems. Deliberately tampering with the system, he deleted budgets and salary information and caused five of the eight network servers to crash. When Federal Bureau of Investigation (FBI) agents caught the perpetrator, his home was filled with hacking tools, proprietary information from Forbes, Inc., and other incriminating material.

In February 1998, hackers from Israel and Northern California attacked the U.S. Department of Defense (DoD). In a carefully organized attack that exploited buffer

overflow, these hackers systematically perpetrated a DoS attack that lasted over a week on 11 DoD sites.

In March 1998, all across the United States, system administrators found their Windows NT servers under apparently automated attack. Systems crashed repeatedly until they were updated to the latest patches from Microsoft. There appeared to be no file damage from the attacks, which lasted more than a day. Sites affected included several NASA and other military sites, as well as several University of California and other college campuses.

Yet another mail storm erupted in May 1998 when an Australian official set autoreply on his email package while he was away. Unfortunately, he inadvertently set his destination for these largely useless messages to be all 2,000 users on his network—and requested auto confirmation of delivery of each autoreply, which generated yet another autoreply, and so on *ad infinitum*. Within four hours, his endless positive-feedback loop had generated 150,000 messages before his autoreply was shut down. The ripples lasted for days, with the perpetrator saddled with 48,000 messages in his IN basket and a stream of 1,500 a day pouring in. This was another case of an inadvertent DoS attack.

In January 1999, someone launched a sustained DoS attack on Ozemail, an important Australian Internet service provider. Email service was disrupted for users in Sydney.

In March 1999, the Melissa email-enabled virus/worm swept the world in a few days as it sent copies of itself to the first 50 addresses in victims' email address books. Because of this high replication rate, the virus spread faster than any previous virus in history. On many corporate systems, the rapid rate of internal replication saturated email servers with outbound automated junk email. Initial estimates were in the range of 100,000 downed systems. Antivirus companies rallied immediately, and updates for all the standard products were available within hours of the first notices from the CERT/CC. The Melissa macro virus was quickly followed by the PAPA MS-Excel macro virus with similar properties but that, in addition, launched DoS attacks on two specific IP addresses.

And after all of the publicity dealing with the Y2 K problems prior to the turn of the millennium, the ILOVEYOU, or Love Letter, worm was released in May 2000. The Love Letter worm was in a Visual BASIC attachment to emails with the subject line "ILOVEYOU." This worm used a so-called double extension exploit; the attachment was named LOVE-LETTER-FOR-YOU.txt.vbs, but because the default Windows setting is to not display the extension, many (or most) recipients thought that this was a harmless text file. The worm modified the infected system's registry to autostart the program on system startup, replaced files with certain extensions with copies of itself, and sent copies of itself to the first 50 entries in the Outlook address book.

More recent attacks are discussed in the sections describing modern DoS tools.

18.3 COSTS OF DENIAL-OF-SERVICE ATTACKS. What are the effects of these DoS attacks in terms of productivity and actual financial costs? It is difficult to place an exact monetary figure on DoS attacks. DoS attacks can interrupt critical processes in an organization, and such interruption can be costly. When a company's computer network is inaccessible to legitimate users and they cannot conduct their normal business, productivity is lowered. The negative effect is bound to carry over to the financial aspects of the business. However, putting exact figures to these effects is uncertain at best, and the estimates are widely disputed, even among security and business experts. In addition, many companies do not comment on the exact losses they suffer because they fear that the negative publicity will decrease their market share. This latter point is significant: In an early 1990s study of Wall Street firms, some of

the companies suggested that if they were to be without their network for two to three days, they might never reopen their doors.

In the case of the Christmas Tree worm, it took IBM several days to clean up its network and resulted in the loss of millions of dollars, both for cleansing the system and for lost business because of lost connectivity and related productivity. Additionally, there was the embarrassment suffered by IBM, a noted technology company. The individual who launched the worm was identified and denied access to any computer account, while IBM had to write a letter of apology to the European Academic and Research Network (EARN) administrators.

In 1988, at the time of the Morris Worm, the Internet consisted of 5,000 to 10,000 hosts, primarily at research and academic institutions. As a result, although the Morris Worm succeeded in bringing many sites to a halt and gained worldwide notoriety, the financial and productivity impact on the commercial world was minimal. A similar incident today would wreak havoc and cost tens or hundreds of millions of dollars in losses.

By 1996, however, commercial reliance on the Internet was already becoming a matter of course. Working around the clock, the management at PANIX and Cisco, the ISP's router vendor, kept the service provider up and running, but the network received 210 fraudulent requests per minute. Although the systems did not crash, thousands of subscribers were unable to receive their email messages. Other sites were attacked in the same time frame as PANIX, including Voters Telecommunication Watch. No one took responsibility for these attacks and it has been widely assumed that they were triggered by articles on SYN DoS attacks (see the description below) that had recently appeared in *2600 Magazine* and *Phrack*, journals that cater to hackers.

According to Forbes, Inc., the losses suffered by the firm because of the DoS attack perpetrated by the disgruntled former employee exceeded $100,000. Could it have been prevented? According to the firm, it is highly unlikely, since Forbes had no reason to suspect the individual was maintaining either the firm's confidential and sensitive material at home or that he was thinking of hacking into the computer system and deliberately doing damage. Although the firm had security on its systems, the perpetrator used the password of a legitimate, authorized user.

The DoS attack launched against the Department of Defense computers in 1998 proved that attackers could deny access to vital military information. In this particular instance, the attack was directed at unclassified machines that had only administrative and accounting records, but it was a blow to the confidence of the Department of Defense. Tying up the computers for over a week presumably reduced productivity, but the government would not comment on the actual cost from loss of machine time and personnel productivity.

These cases show that a DoS attack on a computer or network can be devastating to an organization. Important equipment and networks and even an entire organization can be disabled by such attacks.

Early DoS incidents often were described as annoying, frustrating, or a nuisance. However, with increasing sophistication and dependency on networking, it has become difficult to keep a sense of humor about such incidents. Especially in corporations, where the mission is to make a profit for the shareholder, company managers find it increasingly difficult to excuse being incapacitated because of a DoS or DDoS attack.

As these forms of attack become more sophisticated, so must the tools and methods for detecting and fighting them. Current products scan equipment and networks for vulnerabilities, trigger alerts when an abnormality is found, and frequently assist in eliminating the discovered problem.

18.4 TYPES OF DENIAL-OF-SERVICE ATTACKS. DoS attacks, whether accidental or deliberate, result in loss of service; either a host or a server system is rendered inoperable or a network is rendered inaccessible. DoS attacks are launched deliberately by an *intruder* (the preferred term for attacker in this context). Systems and networks that are compromised are referred to as the *victims*. And while DoS attacks can be launched from the intruder's system, they often are launched by an automated process that allows the intruder to start the attack remotely with a few keystrokes. These programs are known as *daemons,* and they are often placed on another system that the hacker has already compromised.

There are four basic types or categories of DoS attack:

1. **Saturation**. This type of attack seeks to deprive computers and networks of scarce, limited, or nonrenewable resources that are essential in order for the computers or networks to operate. Resources of this type include CPU time, disk space, memory, data structures, network bandwidth, access to other networks and computers, and environmental resources such as cool air and power.
2. **Misconfiguration**. This type of attack destroys or alters configuration information in host systems, servers, or routers. Because poor or improperly configured computers may fail to operate or operate inadequately, this type of attack can be very severe.
3. **Destruction**. This type of attack results in network components being physically destroyed or altered. To guard against this type of attack, it is necessary to have good physical security to safeguard the computers and other network components. This chapter does not deal with physical damage, but Chapters 22 and 23 of this *Handbook* do so in some detail.
4. **Disruption**. This attack interrupts the communications between two devices by altering state information—such as the state of a TCP virtual connection—such that effective data transfer is impossible.

18.5 SPECIFIC DENIAL-OF-SERVICE ATTACKS. The discussion below describes some specific DoS attacks as examples of the general methods that can be employed in a DoS or DDoS. The list is not exhaustive.

18.5.1 Destructive Devices. Destructive devices are programs that accomplish either harassment or destruction of data. There are mixed opinions regarding how severe destructive devices are, but if they threaten a computer's or network's ability to function properly and efficiently, then they may be the instruments of DoS attacks. Viruses, email bombs, and DoS tools all can be considered destructive devices. In fact, viruses and email bombs are known to cause DoS attacks. Some viruses or other malware are known to attack control systems for hardware, such as those employed the Stuxnet and Flame attacks. Viruses and other malicious software and their actions are covered in Chapter 16 of this *Handbook*; nonreplicating or *traditional* DoS and DDoS tools are discussed here.

18.5.2 Email (and Email Subscription) Bombing. Email and email subscription *bombings* were among the first documented DoS attacks. An *email bomb* consists of large numbers of email messages that fill up a victim's electronic mailbox. A huge number of messages can tie up an online connection, slow down mail delivery, and even overload the email server system until the system crashes. Most email bombings

are thought to be deliberate attacks by disgruntled people; specific targets may be victims of someone with a particular grudge. For example, a San Francisco stockbroker received 25,000 messages on September 23, 1996, consisting of the word "Idiot" from a consultant with whom he had had a disagreement. The flood of messages prevented him from using his computer, so in December the victim sued the perpetrator's employer for $25,000 of damages. On occasions in the past, such as with the Christmas Tree worm and the Internet Worm, the DoS is thought to have been accidental.

Email bomb packages automate the process of launching and carrying out an *email bombing* DoS attack. With names like Up Yours, Kaboom, Avalanche, Gatemail, and the Unabomber, these packages can be placed on a network server during a DoS attack and used to attack other systems. Administrators who are aware of these names and others should regularly scan their drives for associated filenames and eliminate them.

To safeguard computers and/or servers, mail filters and exclusionary schemes can automatically filter and reject mail sent from a source address using email bomb packages. Mail filters are available for UNIX, Windows, Macintosh, and Linux systems. Most computer operating systems and most ISPs now offer filtering tools for eliminating unsolicited commercial email and other email. Although perpetrators often disguise their identity and location by using a false address, most filters can be set to screen and eliminate these addresses.

With *email subscription bombing,* also known as *list linking*, a user is subscribed to dozens of mailing lists by the attacker without the user's knowledge. For example, someone calling himself "johnny xchaotic" perpetrated one of the earliest subscription-bombing incidents. In August 1996, he claimed the blame for a massive mail-bombing run based on fraudulently subscribing dozens of victims to hundreds of mailing lists. In a rambling and incoherent letter posted on the Internet, he made rude remarks about famous and not-so-famous people whose capacity to receive meaningful email was obliterated by up to thousands of unwanted messages a day. Today, filtering packages have point-and-click mechanisms that provide automatic list linking. A user could, conceivably, start to receive hundreds or thousands of mail messages per day if linked to just 50 to 100 lists. Once linked to the various lists, the victim has to manually unsubscribe from each individual mailing list. If an attack takes place while the victim is away and without access to email, a user could have a backlog of thousands of messages by the time he or she returns.

List server software should never accept subscriptions without sending a request for confirmation to the supposed subscriber; in the current environment, most list servers require confirmation from the supposed subscriber before initiating regular mailings. The default for the confirmation messages is to ignore the announcement if it is fraudulent. However, even this safety mechanism can generate a wave of many single email messages if a mail bomber abuses the list servers. As a result, many sites have incorporated Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA) images to interfere with bots that attempt to subscribe victims to their mailings.

Speaking of being away, vacation messages and return receipts are another way in which individuals can inadvertently start an *email storm*. Many users set their email clients to automatically request a return receipt of all messages sent. Then the users go on vacation and set up an auto-reply vacation message. When they get a message, the client sends back the vacation message and also requests a receipt. The returned receipts, in turn, generate more auto-reply vacation messages.

Another variant on this feedback loop occurs when an employee goes on vacation and forwards all email to an external ISP that has a local access number in the locale

of the vacation. If the employee decides not to check the mail while away, either due to a high local access fee or so as not to interfere with the vacation, the ISP mailbox will fill up with forwarded messages. If the mailbox fills up, the ISP will send a bounce message back to the corporate server—which then forwards the bounce message back to the ISP, which generates yet another bounce message. Eventually, even the corporate mail server will fill up with a single individual's messages, causing an email DoS.

18.5.3 Buffer Overflow. *Buffer overflow* attacks can be insidious and damaging. It is possible to send an input string to a target program that contains actual code and is long enough to overflow the memory space or input buffer. Sometimes this surreptitious code is placed on the process stack (the area in a computer's memory where the operating system keeps track of the program's input and related code used for processing the inputs), and the code then is processed. An overflow can occur when the input data overflows its buffer space and flows into the stack, where it overwrites the previous data and return address. If the program is written so that the stack address points to the malicious code located in the return buffer, the code executes with the original program's privileges. Buffer overflow is the result of poor programming, where the programmer does not check the size of the input compared to the input buffer. Although the bad-programming basis of buffer overflows should have been eradicated by now, new buffer overflow attacks pop up monthly. As of January 2013, the National Vulnerability Database[1] included 6,273 buffer overflows out of a total of 64,398 vulnerabilities—about 10 percent. This percentage has been dropping over the years; in the period from January 2010 through December 2012, 1,281 of the 14,510 records were for buffer overflows—only 9 percent; in contrast, of the 39,822 records for vulnerabilities entered before 2010, 5,047 (13 percent) involved buffer overflows.

Not all buffer overflows allow the user to insert executable code. DoS attacks such as the *ping of Death* merely attach a data block that is larger than allowed by the IP protocol (i.e., greater than 65,536 bytes). Because the packets are broken into fragments for transmission, they manage to get through the network and, probably, the router and firewall. Once reassembled at the target, however, the packets cause the IP kernel's buffer to overflow and, if not managed properly, the system crashes.

Another example is an old flaw in Microsoft's Internet Information Services[2] (IIS) that could be exploited to allow the Web service to be halted. To do this, an attacker would request a document with a very long URL from an IIS-based Website (and how does one identify an IIS site? If a Website use pages with the .htm or .asp extensions, it is a good guess that the site is running IIS). Upon receipt of the request, an access violation occurred and the server would halt. Although Microsoft issued a patch for this vulnerability, successful attacks continued to take place for years.

18.5.4 Bandwidth Consumption. *Bandwidth consumption* involves generating a large number of packets directed to the network under attack. Such attacks can take place on a local network or can be perpetrated remotely. If the attacker has, or can access, greater bandwidth than the victim has available, the attacker can flood the victim's network connection. Such saturation can happen to both high-speed and low-speed network connections. Although any type of packet may be used, the most common are Internet Control Message Protocol (ICMP) *echo* messages (generated by *pinging*). By engaging multiple sites to flood the victim's network connection, attackers can amplify the DoS attack. To do this successfully, the attackers convince the amplifying system to send traffic to the victim's network. Tracing the intruder who

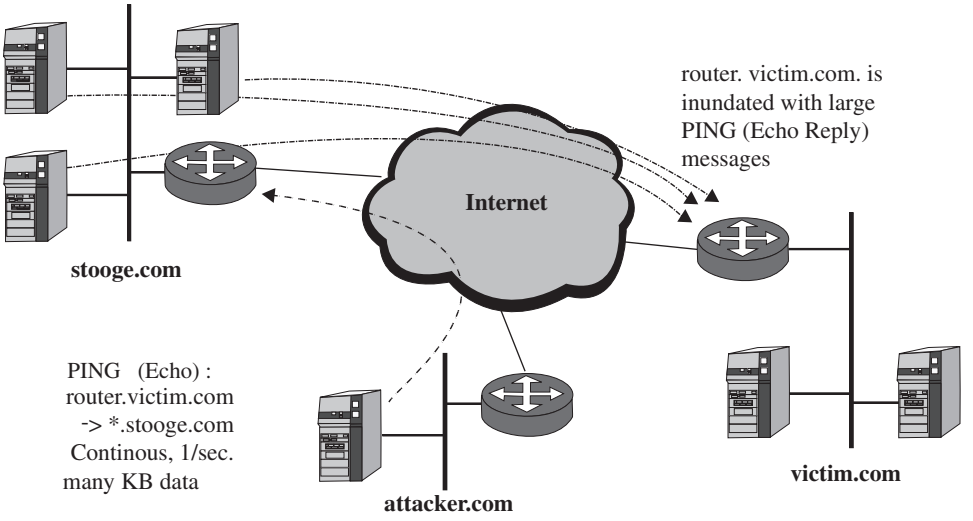


Exhibit 18.1 smurf DoS Attack

perpetrates a bandwidth consumption attack can be difficult since attackers can spoof their source addresses.

One very common bandwidth consumption attack is the *smurf attack*. This attack is particularly interesting (and clever) since it uses tools that reside on every IP system and employs a third-party site without actually having to take control of any system anywhere. As Exhibit 18.1 shows, the smurf attack starts when the intruder (at *attacker.com*) sends a continuous string of very large *ping* messages to the IP broadcast address (* in the exhibit, representing the all-ones, or broadcast, IP address) of the unsuspecting third party, *stooge.com*. The intruder spoofs the source IP address of the *ping* message so that it appears that these messages come from, say, the router at the target network, *router.victim.com*. If the intruder sends a single 10,000-byte *ping* message to the broadcast address of an intermediate site with 50 hosts, for example, the responses will consume 4 megabits (Mb).[3] Even if *victim.com* has a T1 line (with a bandwidth of 1.544 Mb per second), the attacker can swamp a victim's line merely by sending a single large *ping* per second to the right stooge site. The intermediate site is called, for obvious reasons, the *amplifying network*; note that the attacker does not need to compromise any systems there. The ratio of originally transmitted packets to the number of systems that respond is known as the *amplification ratio*.

A variant of the smurf attack is called *fraggle*. In this variant, the attackers send spoofed User Datagram Protocol (UDP) packets instead of *echo* messages to the broadcast address of the amplifying network. Each system on the amplifying network that has the specific broadcast address port enabled will create large amount of traffic by responding to the victim's host; if the port is not enabled, the system on the amplifying network will generate ICMP *host unreachable* messages to the victim's host. In either case, the victim's bandwidth is consumed.

Kernel panic attacks are not due to programming error, per se, but to a hole in a program's logic. As an example, the Intel Pentium chip that could not correctly divide two particular legal inputs had a programming flaw. An IP kernel that fails when receiving a packet that should never occur in the first place is, indeed, a gap in the program's logic but not the same as failing to handle legal input. In this case, the program did not know how to fail gracefully when it received unexpected input.

A specific example of kernel panic occurs in Linux kernel v.2.2.0 when a program usually used for printing shared library dependencies is used instead to print some core files. Under certain circumstances, *munmap()*, a function call used to map and unmap devices into memory, overwrites critical areas of kernel memory and causes the system to panic and reboot.

And there are other examples of these kinds of attacks. A *land attack* occurs when a spoofed packet is sent to a target host where the TCP source and destination ports are set to the same value, and the IP source and destination addresses are set to the same value. Since this is confusing to the host operating system, it results in 100 percent CPU utilization and then a halt. Land attacks have been directed at just about all operating systems.

Teardrop attacks are also the result of behavior when receiving impossible packets. If an IP packet is too large for a particular network to handle, the packet will be fragmented into smaller pieces. Information in the IP packet header tells the destination host how to reassemble the packet. In a Teardrop attack, the attacker deliberately crafts IP packets that appear to overlap when reassembled. This, too, can cause the host to crash. Teardrop attacks have been targeted against Microsoft operating systems and all variants of UNIX.

18.5.5 Routing and Domain Name System Attacks. *Routing* and *Domain Name System (DNS) attacks* are clever attacks that are achieved repeatedly. By tampering with the DNS, a site's domain name resolves to the IP address of any other site that the attacker wishes. In August 1999, for example, a disgruntled engineer for PairGain Technologies provided a direct link to a bogus but authentic-looking Bloomberg News Service page he had created. The engineer, who owned PairGain shares, posted false information about PairGain's supposed acquisition by an Israeli company, pumping up the price of PairGain stock worldwide and creating havoc in the market through his page-jacking exploits. Although a nontraditional DoS attack, these actions did deny users access to the site that they desired, with serious consequences.

Another example occurred in July 1997, when Eugene Kashpureff filed fraudulent information with InterNIC for its DNS updates for that month, forcing domain name servers around the globe to recognize temporary and unauthorized Internet addresses ending in .xxx, .mall, .nic, and .per. A few weeks later, Kashpureff inserted false information that forced people trying to access the Website of Network Solutions Inc. to end up at the AlterNIC site, which he owned. He was arrested later that year on charges of wire fraud.

In another example—this from 2000—of a routing and DNS attack, RSA Security, Inc., after announcing that it had developed a method to combat Website hackers, found that users were unwittingly being rerouted to a counterfeit RSA Website. The fraudulent site looked exactly like the original RSA Website but made fun of the fact that the hacker had managed to achieve his DoS goal.

As far back as 1997, weaknesses were found in and documented about BIND implementations in versions preceding 4.9.5+P1. The earlier versions would cache bogus DNS information when DNS recursion was enabled. The vulnerability enabled attackers to exploit the process of mapping IP addresses to hostnames in what is known as *pointer (PTR) record spoofing*. This type of spoofing provides the potential for a DNS DoS attack.

DNS attacks are still seen widely today. *Phishing* is a form of social engineering attack whereby users are directed to bogus, but authentic-looking, bank or credit card company Websites and enticed to enter personal information used for identity theft.

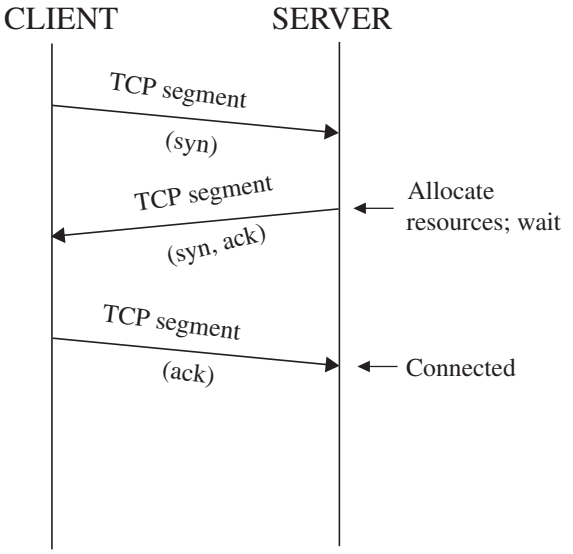


EXHIBIT 18.2 Normal TCP 3-Way Handshake

If users were to look closely at the URL of the site, however, they would observe a suspicious URL. *Pharming* is a variant of phishing that relies on some form of DNS poisoning so that a user going to the actual URL of a bank or credit card company will be redirected to the bogus site. Other variants of phishing have also emerged, such as *spearphishing* (directing a phishing attack at a specific person, job function, or group), *vishing* (a phishing attack via the telecommunications network, such as Voice-over-IP [VoIP]), and *smishing* (phishing via Short Message Service [SMS], or text, messages). For more details of such social engineering attacks, see Chapter 19 in this *Handbook*.

18.5.6 SYN Flooding. *SYN flooding* is a DoS attack that fits into the consumption-of-scarce-resource category. This DoS attack exploits the three-way handshake used by TCP hosts to synchronize the logical connection prior to data exchange. In a normal TCP host-to-host connection, the two hosts exchange three TCP segments prior to exchanging data, as shown in Exhibit 18.2:

1. The client sends a segment to the server with its initial sequence number (ISN). The SYN (synchronization) flag is set in this segment.
2. The server responds by sending a segment containing its ISN and acknowledges the client's ISN. This segment will have both the SYN and ACK (acknowledgment) flags set. At this point, the server allocates resources for the soon-to-be-established connection and waits for the third segment.
3. The client sends a segment acknowledging the server's ISN. This and all subsequent segments until the end of the session will have only the ACK flag set.

A SYN flood takes advantage of the three-way handshake and the fact that a server can have only a finite number of open TCP connections. The attack is launched when an attacker initiates connections for which there will never be a third segment (see Exhibit 18.3). After the server sends the segment in step 2, it waits for a response. Under normal circumstances, the client will respond within a few seconds. The server

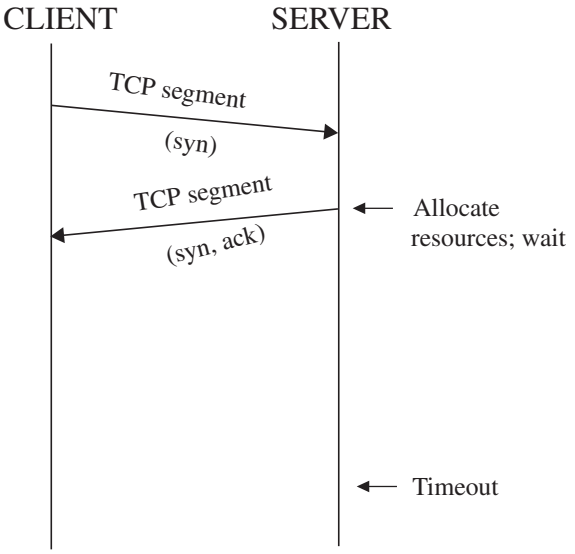


EXHIBIT 18.3 TCP SYN DoS Attack

might wait, say, 10 seconds before timing out and releasing the resources. But suppose the attacker sends hundreds of connection messages per second on a sustained basis. When these bogus connection attempts flood the target faster than they can time out, there will not be any resource left in which to establish a legitimate connection. This is the type of attack that was launched against PANIX in 1996.

18.5.7 Resource Starvation. *Resource starvation* is a catchall category for many other types of DoS attacks that are the result of some scarce resource—be it bandwidth, processing power, disk space—being consumed and exhausted. One example uses a novel UDP DoS attack, where an intruder forges UDP packets and uses them to connect the *echo* service of one machine (UDP port 7) to the character generator (*chargen*) service on another machine (UDP port 19). When this occurs, the two machines consume all available network bandwidth, and all machines on the same networks as the target machines can be affected and starved for resources.

But such attacks can happen locally, as well. Unfortunately, many authorized users carry out *local DoS attacks* that either consume system resources or deny users access. Numerous resource starvation attacks and programming flaws attack specific systems. For instance, by exceeding imposed quotas on disk space, multiuser systems can suffer from a resource starvation attack. Most operating systems limit the amount by which a user can exceed disk quota. Windows 2000 has a twist on this; although individual files can exceed their quota by only a small amount, if *every* file exceeds the quota, a user can consume a lot of extra disk space.

Another variant of this type of attack is called *Slowloris*, which targets Web servers and uses a relatively low amount of bandwidth by the attacker. Using the Slowloris scheme, an attacker sends a large number of connection requests to a target Web server and holds them open as long as possible. Under some circumstances, the affected Web server will exhaust its pool of connection resources on unfulfilled requests, thus denying legitimate requests. Many types of Web servers are susceptible to Slowloris, including several versions of Apache.

A Slowloris-variant is called *R-U-Dead-Yet* (RUDY). RUDY sucks connections resources from a Web server by sending large numbers of POST messages and randomly large Content-Length values in the HTTP message header.

18.5.8 Java, PHP, and ASP Attacks. There are many attacks that target program vulnerabilities in Websites that use Java, PHP (PHP: Hypertext Preprocessor), Active Server Pages (ASP), and other server-side scripting languages. One example is a 2011 DOS attack that resulted in CPU resource exhaustion. In this particular attack, an attacker would send variables to a Web server using the Hypertext Transfer Protocol (HTTP) POST method. By sending tens or hundreds of thousands of specially selected variable names, collisions in the hash table names could be forced, eventually consuming all CPU cycles on handling a single HTTP POST request.

18.5.9 Router Attacks. *Router attacks* have been developed for a wide variety of routers. Because routers are the backbone of the Internet and the gateway through which organizations connect to the Internet, killing a router denies network service for hundreds of machines. Productivity and financial repercussions from a hardware attack can be very severe. Popular targets for these kinds of attacks are routers from Ascend, Cisco, Juniper, Lucent, and 3Com. Unfortunately, many network managers make the attacks easy by employing Telnet or HTTP for remote access and not properly securing the network against remote access by anyone over the Internet.

18.5.10 Other Denial-of-Service Attacks. Even this list of specific DoS attacks is not exhaustive. For example, *Bonk* and *boink* (aka *bonk.c*) are DoS attacks that cause a target Windows 9x/NT system to crash. *Arnudp* (*arnudp100.c*) forges UDP packets to implement a DoS against UDP ports 7 (*echo*), 13 (*daytime*), 19 (*chargen*), and 37 (*time*), services that frequently run on UNIX and Windows NT systems. And *cbcb.c* is a *cancelbot* that destroys existing Usenet news postings that fit certain criteria. (Some argue that *cbcb.c* does not actually carry out a DoS attack; however, once activated, the program denies access to postings for targeted Usenet users.)

An example of network-based attack on a non-network resource would be *billion laughs* (an XML bomb), a DoS attack targeting Extensible Markup Language (XML) parsers. In XML, an *entity* is a piece of code that allows the programmer to define something that can be easily modified or reused in other documents, not unlike a programming object. In *billion laughs*, the attacker creates an XML document containing 10 entities, each defined as being composed of 10 instances of the previous entity ($10^{10} \approx 1$ billion, hence the name). Billion laughs was first introduced as early as 2003 but remained a problem late into at least the first decade of the 2000s.

On a final note, DoS attacks are not limited to networks. A *black fax* attack, for example, is a simple yet effective resource depletion attack; the attacker merely faxes a page containing a large black box to the victim, which causes the victim's fax to use a large amount of black toner. Sending a large number of these faxes will ultimately deplete the toner at the victim's fax and/or prevent legitimate faxes from being delivered. A security expert once glued the ends of a long strip of paper together to create an endless loop in a fax machine to attack a fax-spammer who refused to stop sending him unwanted travel offers. Another attack on faxes uses automated scripts for a computer to dial and redial the target machine with thousands of copies of (usually offensive) faxes; irritated nerds have been known to use this attack on fax-spammers unwise enough to use real telephone numbers for their own fax machines.

18.6 PREVENTING AND RESPONDING TO DENIAL-OF-SERVICE ATTACKS. DoS attacks are best prevented; handling them in real time is very difficult. And the most important way to protect a system is to harden the operating systems:

- Install them with security in mind.
- Monitor sites to be aware of security vulnerabilities.
- Maintain the latest versions of software where possible.
- Install all relevant security patches.

But a large measure of the prevention consists of packet filtering at network routers. Because attackers frequently hide the identity of the machines used to carry out the attacks by falsifying the source address of the network connection, techniques known as *egress filtering* and *ingress filtering* are commonly used as protective measures. As discussed later in this chapter, egress and ingress filtering are methods of preventing packets from leaving or entering the network, respectively, with an invalid source address. Blocking addresses that do not fit the criteria for legitimate source addresses and making certain that all packets leaving an organization's site contain legitimate addresses can thwart many DoS attacks.

Other packet-filtering methods that will help prevent DoS are to block all broadcast messages and most ICMP messages. There is no reason that a site should accept messages being broadcast to all hosts on the site. Furthermore, there is probably no good reason to allow all hosts to respond to *ping* or traceroute messages; in fact, most ICMP messages probably can be blocked.

In some instances, victims have set up response letters triggered to send and resend in large quantities so that they flood the attacker's address. Doing this is generally not a good idea. If these messages are sent to a legitimate address, the attacker may get the message and stop. But the attackers generally spoof the source IP address, so responding back to a possibly bogus IP network is not a good defensive posture because it may harm innocent victims. The best defense will involve the ISP.

In instances where the attacker's service provider can be identified and contacted, the victim can request that the service provider intervene. In these instances, it is usual for the ISPs to take appropriate action to stop the attack and find the perpetrator. However, in instances where a DoS appears to emulate or mimic another form of attack or when it continues for an unusually long period of time, the victim may want to take more aggressive action by contacting CERT/CC, the FBI, and other authorities that have experience with DoS attacks and some jurisdiction if the perpetrators are caught.

Real-time defenses are difficult but possible. Many routers and external intrusion detection systems (IDSs) can detect an attack in real time, such as too many connection requests per unit time from a given IP host or network address. A router might block the connection requests or an IDS might send a pager message to a security administrator.

However, attacks such as smurfs can suck up all of the bandwidth even before the packets get to the target site. Cooperation by ISPs and end-user sites is required to fully combat DoS attacks. This will be addressed further as part of the discussion of responding to DDoS.

18.7 DISTRIBUTED DENIAL-OF-SERVICE ATTACKS. DDoS tools use amplification to augment the power of the attacker. By subverting poorly secured systems

into sending coordinated waves of fraudulent traffic aimed at specific targets, intruders can overwhelm the bandwidth of any given victim.

In a DDoS attack, the attacking packets come from tens or hundreds of addresses rather than just one, as in a standard DoS attack. Any DoS defense that is based on monitoring the volume of packets coming from a single address or single network will fail since the attacks come from all over. Rather than receiving, for example, 1,000 gigantic *pings* per second from an attacking site, the victim might receive one *ping* per second from each of 1,000 attacking sites.

One of the other disconcerting things about DDoS attacks is that the handler can choose the location of the agents. So, for example, a handler could target several North Atlantic Treaty Organization (NATO) sites as victims and employ agents that are all in countries known to be hostile to NATO. The human attacker might be sitting in, say, Canada.

Like DoS attacks, all of the DDoS attacks employ standard TCP/IP messages—but employ them in some nonstandard ways. Common DDoS attacks include *Tribe Flood Network* (TFN), *Trin00*, *Stacheldraht*, and *Trinity*. The sections that follow present some details about these attacks.

18.7.1 Short History of Distributed Denial of Service. Denial-of-service attacks under a number of guises have been around for decades. Distributed DoS attacks are much newer. In late June and early July 1999, groups of hackers installed and tested a DDoS tool called *Trinoo* (see Section 18.7.3.1) to launch medium to large DDoS attacks. Their tests involved over 2,000 compromised systems and targets around the world.

Most of the literature suggests that the first documented large-scale DDoS attack occurred in August 1999, when Trinoo was deployed in at least 227 systems (114 of which were on Internet2) to flood a single University of Minnesota computer; this system was down for more than two days.

On December 28, 1999, CERT/CC issued its Advisory CA-1999-17 (www.cert.org/advisories/CA-1999-17.html) reviewing DDoS.

On February 7, 2000, Yahoo! was the victim of a DDoS during which its Internet portal was inaccessible for three hours. On February 8, Amazon, Buy.com, CNN, and eBay were all hit by DDoS attacks that caused them either to stop functioning completely or to slow down significantly. And on February 9, 2000, E*Trade and ZDNet both suffered DDoS attacks. Analysts estimated that during the three hours Yahoo! was down, it suffered a loss of e-commerce and advertising revenue that amounted to about $500,000. According to bookseller Amazon.com, its widely publicized attack resulted in a loss of $600,000 during the 10 hours it was down. During their DDoS attacks, Buy.com went from 100 percent availability to 9.4 percent, while CNN.com's users went down to below 5 percent of normal volume and Zdnet.com and E*Trade.com were virtually unreachable. Schwab.com, the online venue of the discount broker Charles Schwab, was also hit but refused to give out exact figures for losses. One can only assume that to a company that does $2 billion weekly in online trades, the downtime loss was huge. Another type of damage caused indirectly by the DDoS was the decline in stock values in the 10 days following the attacks: eBay suffered a 24 percent decline, Yahoo! dropped 15 percent, and Buy.com dropped 44 percent.

These attacks were orchestrated by 15-year-old Michael Calce, a youngster living in the west end of Montreal island in Quebec whose pseudonym was *Mafiaboy*. He eventually pled guilty to 56 charges of computer crime and served eight months of juvenile detention.[4]

These types of DDoS attacks have continued since the summer of 1999. One of the best-known incidents was a series of DDoS attacks again Steve Gibson's GRC.com Website in May 2001. The attacker was a 13-year-old using an Internet Relay Chat (IRC) bot, automated programs that exploit systems using IRC clients to become DDoS zombies.

The DDoS attack that had the highest potentially devastating impact, however, occurred on October 21, 2002, when all of the top-level DNS root servers were subjected to a sustained attack by thousands of zombies. Nine of the 13 DNS root servers were knocked off the Internet; the remaining four were able to keep operating during the attack. All of the major ISPs and many large private networks maintain their own DNS systems, although most servers ultimately rely on the root servers to find noncached DNS entries. The attack lasted for just an hour or two; had it continued for much longer, the remaining servers probably would have been overwhelmed, effectively blocking DNS host name/address translation.

A disturbing and growing trend is the use of DDoS as an extortion tool. An increasing number of criminals are using DDoS tools as a way to threaten an attack rather than actually disrupting a target organization's network. Although several providers of network, security, and consulting services claim that they and many of their customers have received such extortion demands, few are public about naming the targets—many of which are acceding to the threats and paying the blackmail.[5] Most experts agree that the extortion demands should not be met; doing so only encourages the criminal behavior. If such a threat is received, the organization's ISP and law enforcement authorities should be contacted immediately.

18.7.2 Distributed Denial-of-Service Terminology and Overview. To describe and understand DDoS attacks, it is important to understand the terminology that is used to describe the attacks and the tools. Although the industry has more or less settled on some common terms, that consensus did not come about until well after many DoS/DDoS attacks had already appeared in the hacker and mainstream literature. Early descriptions of DDoS tools used a jumble of terms to describe the various roles of the systems involved in the attack. At the CERT/CC Distributed System Intruder Tools Workshop held in November 1999, some standard terminology was introduced and those terms are used in the paragraphs below. To align those terms and the terms used by the hacker literature as well as early descriptions, here are some synonyms:

Intruder—also called the *attacker* or *client.*

Master—also called the *handler.*

Daemon—also called an agent, bcast (broadcast) program, or zombie.

Victim—also called the *target.*

DoS/DDoS attacks actually have two victims: the ultimate target and the intermediate system(s) that were exploited and loaded with daemon software. In this chapter, the focus is on the end-of-the-line DoS/DDoS victim.

DDoS attacks always involve a number of systems. A typical DDoS attack scenario might follow roughly these three steps:

1. The *intruder* finds one or more systems on the Internet that can be compromised and exploited (see Exhibit 18.4). This is generally accomplished using a stolen

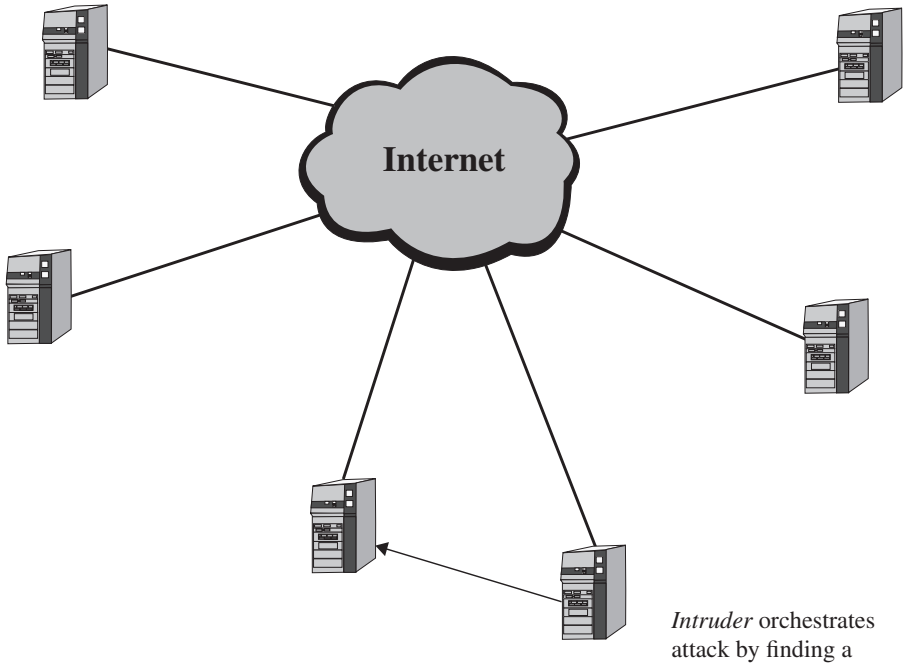


EXHIBIT 18.4 DDoS Phase 1

account on a system with a large number of users or inattentive administrators, preferably with a high-bandwidth connection to the Internet. (Many such systems can be found on college and university campuses.)

2. The compromised system is loaded with any number of hacking and cracking tools, such as scanners, exploit tools, operating system detectors, rootkits, and DoS/DDoS programs. This system becomes the DDoS *master*. The master software allows it to find a number of other systems that can themselves be compromised and exploited. The attacker scans large ranges of IP network address blocks to find systems running services known to have security vulnerabilities. This *initial mass-intrusion phase* employs automated tools to remotely compromise several hundred to several thousand hosts, and installs DDoS agents on those systems. The automated tools to perform this compromise are not part of the DDoS toolkit but are exchanged within groups of criminal hackers. These compromised systems are the initial victims of the DDoS attack. These subsequently exploited systems will be loaded with the DDoS *daemons* that carry out the actual attack (see Exhibit 18.5).
3. The intruder maintains a list of *owned systems* (sometimes spelled *pwned* by hackers), the compromised systems with the DDoS daemon. The actual *denial-of-service attack phase* occurs when the attacker runs a program at the master system that communicates with the DDoS daemons to launch the attack. Here is where the intended DDoS victim comes into the scenario (see Exhibit 18.6).

Communication between the master and daemons can be obscured so that it becomes difficult to locate the master computer. Although some evidence may exist on one or more machines in the DDoS network regarding the location of the master, the daemons

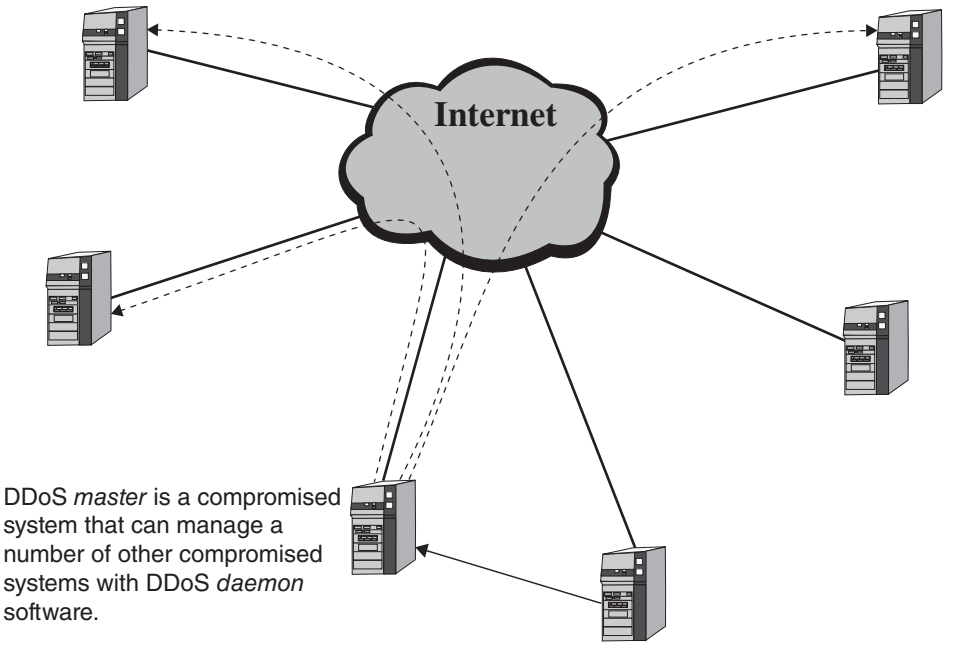


Exhibit 18.5 DDoS Phase 2

normally are automated so that it is not necessary for an ongoing dialog to take place between the master and the rest of the DDoS network. In fact, typically techniques are employed to deliberately camouflage the identity and location of the master within the DDoS network. These techniques make it difficult to analyze an attack in progress and difficult to block attacking traffic and trace it back to its source.

In most cases, the system administrators of the infected systems do not even know that the daemons have been put in place. Even if they do find and eradicate the DDoS

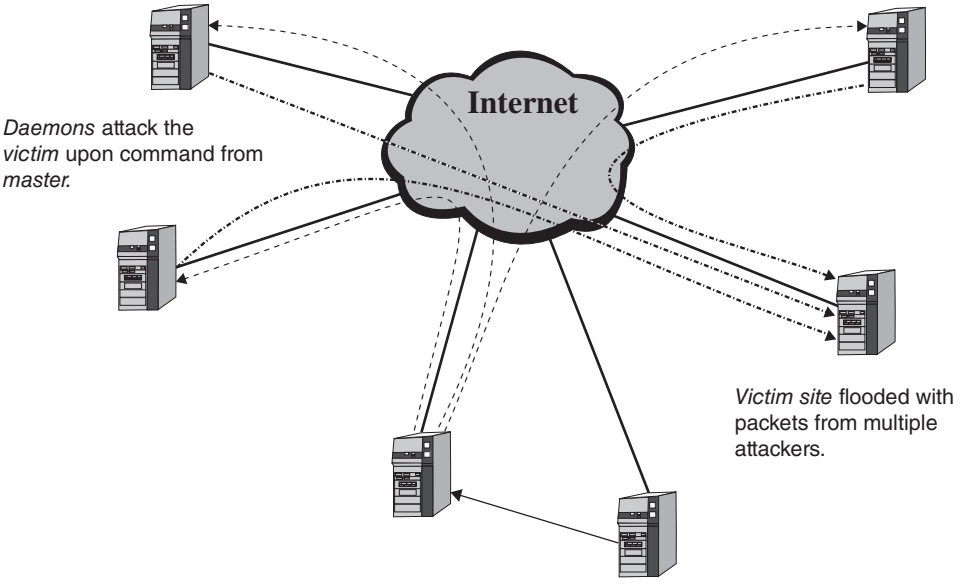


Exhibit 18.6 DDoS Phase 3

software, they cannot help anyone determine where else the software may have been placed. Popular systems to exploit are a site's Web, email, name, or other servers, since these systems are likely to have a large number of open ports, a large amount of traffic, and are unlikely to be quickly pulled off-line even if an attack can be traced to them.

18.7.3 Distributed Denial-of-Service Tool Descriptions. This section provides some details on how some of the major DDoS tools work.

18.7.3.1 Trinoo (Trin00). *Trinoo* or *Trin00* was the first known DDoS tool, appearing in the summer of 1999. The typical installation of Trinoo is similar to the scenario painted above where an attacker plants handler software on a system and the handler, in turn, loads the attack software on the agents. Trinoo is a distributed SYN DoS attack.

Trinoo uses a number of TCP and UDP ports:

- Masters listen on TCP port 27665 for attacker-to-master communication.
- Daemons listen on UDP port 27444 for master-to-daemon communication.
- Masters listen on UDP port 31335 for daemon-to-master communication.

These are default port numbers, and variants used other ports. The human attacker can control a Trinoo master (the handler) remotely via a connection to TCP port 27665. After connecting, the attacker gives the expected password, *betaalmostdone*.

The Trinoo master program is typically named *master.c* and the daemon is *ns.c*. Communication between the Trinoo master (handler) and daemons (agents) is via UDP. Master-to-daemon communications employ UDP datagrams on port 27444. All commands contain a password, the default being *l44adsl*. All valid commands contain the substring *l44*.

Communication from the Trinoo daemons to the master use UDP datagrams on port 31335. When the daemon starts, it sends a message to the master containing the string **HELLO**. The Trinoo *keep alive* function is accomplished by an exchange between the master and daemon: The master sends a Trinoo *mping* command, which sends the string *png* to a daemon; the daemon responds by sending the string *PONG* to the master.

The passwords are there to prevent system administrators from being able to take control of the masters and daemons that form the Trinoo network. Other default passwords in the initial attacks were *gOrave* to start the Trinoo master server and *killme* to control the master's *mdie* command to kill the Trinoo processes. Like the port numbers, the passwords can be changed easily by the attackers.

Intrusion detection software or system management routine analysis can look for a number of things that might indicate the presence of Trinoo:

- A system listening on UDP port 27444 could be a Trinoo daemon.
  - Trinoo daemon communication will contain the string *l44*.
  - The SYN flood mechanism picks the destination port using a random number generator function.
  - A Trinoo daemon will send the string *PONG* if it receives a *png* command.
- A system listening on TCP port 27665 could be a Trinoo master.

- A system listening on UDP port 27444 could be a Trinoo master.
  - UDP packets will contain the string *l44adsl.*

18.7.3.2 Tribe Flood Network. The *Tribe Flood Network* (TFN) appeared after Trinoo. TFN runs primarily on compromised UNIX systems exploited using buffer overrun bugs in the remote procedure call (RPC) service. TFN client and daemon programs implement a DDoS network capable of employing a number of attacks, such as ICMP flood, SYN flood, UDP flood, and smurf-style attacks.

TFN is noticeably different from Trinoo in that all communication between the client (attacker), handlers, and agents uses *ICMP echo* and *echo reply* packets. Communication from the TFN client to daemons is accomplished via *ICMP echo reply* packets. The absence of TCP and UDP traffic sometimes makes these packets difficult to detect because many protocol monitoring tools are not configured to capture and display the ICMP traffic.

Remote control of the TFN network is accomplished by executing a program on the client system. The program also can be executed at the handler system by the client via some host-to-host connection method, such as connecting to an exploited TCP port or using a UDP- or ICMP-based remote shell. The program must be supplied:

- The IP address list of hosts that are ready to carry out the flood attack
- The type of attack to be launched
- The IP address list of target hosts
- The port number for SYN attack

No password protection is associated with TFN. Each command to the daemons is sent in the Identifier field of the ICMP packet; values *345*, *890*, and *901* start the *SYN*, *UDP*, and *ICMP* flood attacks, respectively. The Sequence Number field in the *echo reply* message is always set to *0x0000*, which make it look like the response to the initial *echo* packet sent out by the *ping* command.

The TFN client program typically is named *tribe.c* and the daemon is *td.c*.

18.7.3.3 Stacheldraht. Stacheldraht (German for *barbed wire*) is a DDoS tool that appeared in August 1999 and combines features of Trinoo and TFN. It also contains some advanced features, such as encrypted attacker–master communication and automated agent updates.

Stacheldraht uses a Trinoo-like client/server architecture. The handler listens on TCP port 16660 for client (intruder) commands, and the agents listen on TCP port 65000 for commands from the handler. Agent responses to the handler employ *ICMP echo reply* messages. The possible attacks are similar to those of TFN; namely, ICMP flood, SYN flood, UDP flood, and smurf attacks.

Trinoo and TFN exchange commands in plaintext. Trinoo, being TCP-based, is also subject to common TCP attacks, such as session hijacking. Stacheldraht addresses these deficiencies by employing an encrypting *telnet alike* client. (*Telnet alike* is a Stacheldraht term.) The client uses secret-key cryptography.

The Stacheldraht network comprises a number of programs. The attacker uses an encrypting client called *telnetc/client.c* to control to one or more handlers. The handler program is called *mserv.c,* and each handler can control up to 1,000 agents. The agent software, *leaf/td.c,* coordinates the attack against one or more victims upon command from the handler.

18.7.3.4 TFN2 K. *Tribe Flood Network 2 K* (TFN2 K) was released in December 1999 and targets UNIX and Windows NT servers. TFN2 K is a complex variant of the original TFN with features designed specifically to:

- Make TFN2 K traffic difficult to recognize and filter
- Remotely execute commands
- Hide the true source of the attack using IP address spoofing
- Transport TFN2 K traffic over multiple transport protocols including UDP, TCP, and ICMP
- Confuse attempts to locate other nodes in a TFN2 K network by sending "decoy" packets

TFN2 K, like TFN, can consume all of a system's bandwidth by flooding the victim machine with data. But TFN2 K, unlike TFN, also includes attacks designed to crash or introduce instabilities in systems by sending malformed or invalid packets, such as those found in the Teardrop and Land attacks.

TFN2 K uses a client server architecture in which a single client issues commands simultaneously to a set of TFN2 K agents. The agents then conduct the DoS attacks against the victim(s). The agent software is installed in a machine that already has been compromised by the attacker.

18.7.3.5 Shaft. Trinoo, TFN/TFN2 K, and Stacheldraht were the first, and possibly best studied, DDoS tools. Other tools followed, of course, and have become increasingly more complex, but these early tools remain available and some systems remain vulnerable to these forms of attack.

In November 1999, for example, the *Shaft* DDoS tool became available. A Shaft network looks conceptually similar to a Trinoo network with client-managing handler programs (*shaftmaster*) that, in turn, manage agent programs (*shaftnode*). Like Trinoo, handler-agent communication uses UDP, with the handler(s) listening on port 20433 and the agent(s) listening on port 18753. The client communicates with the handler by telnetting to TCP port 20432. The attack itself is a packet flooding attack, and the client controls the size of the flooding packets and duration of the attack. One signature of Shaft is that the sequence number for all TCP packets is always 0×28374839.

18.7.3.6 HTTP Apache Attack. In August 2000, a DDoS attack against Apache Web servers was first detected. The attack took advantage of a vulnerability whereby a URL sent to an Apache Web server containing thousands of forward slashes (/) would put the server into a state that would consume enormous CPU time. This particular attack was launched by over 500 compromised Windows computers and would, presumably, succeed against Apache Web servers prior to version 1.2.5.

18.7.3.7 Trinity. In September 2000, a new DDoS tool called *Trinity* was reported. Trinity is capable of launching several types of flooding attacks on a victim site, including ACK, fragment, RST, SYN, UDP, and other floods. Trinity agent software must be placed on Linux systems compromised by a buffer overflow vulnerability. The agent binary code typically is found in */usr/lib/idle.so*. Communication from the handler or intruder to the agent, however, is accomplished via Internet Relay Chat (IRC) or America Online's ICQ instant messaging software. Whereas the attacker has

to keep track of the IP addresses of compromised systems with Trinoo and TFN, all of the Trinity agents report back to the attacker by appearing in the same chat room. The original reports were that the Trinity agent communicated over an IRC channel called *#b3eblebr0x*; other IRC channels presumably are being used for DDoS, as well. IRC uses TCP ports 6665 to −6669, and Trinity appears to use port 6667. In addition, a binary called */var/spool/uucp/uucico* is a backdoor program that listens on TCP port 33270 for connections; an attacker connecting on that port and providing the password !@# will achieve rootshell on the affected system.

18.7.3.8 SubSeven. Zombie software is not always distributed by an attacker exploiting a vulnerability of an exposed system. Indeed, very often the user is the culprit. Trojan horses are often the mechanism for distributing the zombie code. The *SubSeven* software, for example, is a backdoor virus. SubSeven often gets on a user's system because it is distributed within programs available via Usenet and other Internet sites, such as some game or pornography programs (e.g., SexxxyMovie.mpeg.exe). Potential attackers frequently scan computer systems today, particularly residential systems connected to the Internet via DSL or cable modem, for the presence of SubSeven, which provides a potential backdoor into users' systems; system administrators also are learning to scan for this dangerous program on their own systems.

18.7.3.9 Mydoom. The description above of some of the first DDoS tools is meant to show the early baseline from which later tools evolved, many of which are simply a variant of the early tools. *Mydoom* (aka W32.MyDoom@mm and other aliases) is an example of a DoS attack via an email worm, first launched in 2004. Mydoom appears as an email with some sort of error-type message in the subject line (e.g., "Error" or "Mail System Failure"). Interestingly, the email could be prepared in several different languages, including English, French, and German. The email contained an attachment that, if executed, forwarded the worm to addresses found in a variety of files on the local system, including the email address book. The worm could also be copied to shared folders of a user's peer-to-peer networking sites. Early versions of Mydoom contained a payload that would launch a DoS attack against the SCO Group on February 1, 2004; the SCO Group was unpopular because they asserted, in a $1 billion lawsuit against IBM, that several Linux distributions were violating SCO's UNIX intellectual assets. Other variants of Mydoom were subsequently released, including a version employed in attacks on South Korea and the United States in July 2009 that included the launching of a botnet.

18.7.3.10 LOIC, HOIC, and HULK. Another variant of these types of attacks is the *Low Orbit Ion Cannon* (LOIC). Designed to be used as a network stress test, it can also be deployed as a DoS or DDoS tool. LOIC works, basically, by flooding a target with TCP and UDP packets in order to learn whether the routers can handle the load and how they will respond—or to usurp all of the bandwidth. Originally written in C#, a JavaScript variant (LS LOIC) and Web version (Low Orbit Web Cannon) have also appeared. LOIC has been used extensively by the hacktivist collective *Anonymous*.

The *High Orbit Ion Cannon* (HOIC) is a modification of LOIC. It has an easy-to-use graphical user interface that allows attackers to specify targets and *booster* scripts to target multiple pages on the targeted server, much like a shotgun instead of a rifle. The *intensity* setting defines the number of requests per second (2/sec for *low* and 8/sec for *high*).

The *HTTP Unbearable Load King* (HULK) is a variant of other tools that bombard a Web server with a overwhelming number of packets. Most prior tools, however, send TCP SYN requests or other predictable packets, which allows the firewall and/or server to detect the attack and mount a defense. HULK generates a unique, nonpredictable set of requests designed to frustrate defenses based on pattern recognition and packet filtering.

18.8 DENIAL-OF-SERVICE USING EXPLOITABLE SOFTWARE. The tools just discussed employ a common DoS approach: An attacker exploits a vulnerability of a potential victim and uses that system to launch attacks on the intended victim. Later rounds of DDoS attacks, however, use code that is commonly available and that has known vulnerabilities. And all too often, the vulnerabilities are exploited after a patch to fix them has been released but ignored by some system administrators. A case in point is the flurry of such releases in mid-2001.

In May 2001, a buffer overflow exploit was discovered in the Microsoft IIS Indexing Service. In mid-June, Microsoft released a security bulletin warning that administrative scripts (.ida files) and Internet data queries (.idq files) did not do proper bounds checking. As it happens, what seems like the vast majority of IIS servers did not get the patch, and, in essence, every unpatched IIS server became a DDoS zombie.

18.8.1 Code Red. In July 2001, eEye Digital Security and several other security organizations around the Internet saw an alarming number of TCP port 80 scans on the Internet. What they eventually discovered was what became known as the *Code Red Worm*.

Code Red had three distinct phases. The *propagation phase* occurred during the first 19 days of the month. During this phase, the attacking system scanned target systems on TCP port 80 and sent a specially crafted HTTP GET request that exploited the IIS buffer overflow (even if the Index Service is not running). A common log entry might appear as:

```
211.5.255.44 - - [16/Aug/2001:11:30:49 -0400] “GET
/default.ida?NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN%u9090%u685
8%ucbd3%u7801u9090%u6858%ucbd3%u7801%u9090%u6858%ucbd3%u7801%u
9090%u9090%u8190%u00 c3%u0003%u8b00%u531b%u53ff%u0078%u0000%u00
=a HTTP/1.0” 400 357
```

If the exploit was successful, the worm ran in RAM of the infected server and spawned 99 new threads to attack a quasi-random set of IP addresses. If the exploited server’s native language was English, the server’s Web page was defaced with a message that said “Welcome to http://www.worm.com! Hacked by Chinese!” This message would stay up for 10 hours and then disappear.

The *flood phase* occurred on days 20 to 27 of the month. This is when the attack really happened; every day between 8:00 and 11:59 P.M. UTC, the compromised servers sent 100 KB packets to the IP address 198.137.240.91, which formerly was assigned to www.whitehouse.gov. (Once the actions of Code Red were discovered, the IP address of www.whitehouse.gov was changed from the targeted address.)

Days 28 to 31 of the month were the *termination phase*, when the worm became dormant. Code Red was relatively innocuous, compared to what it could have been;

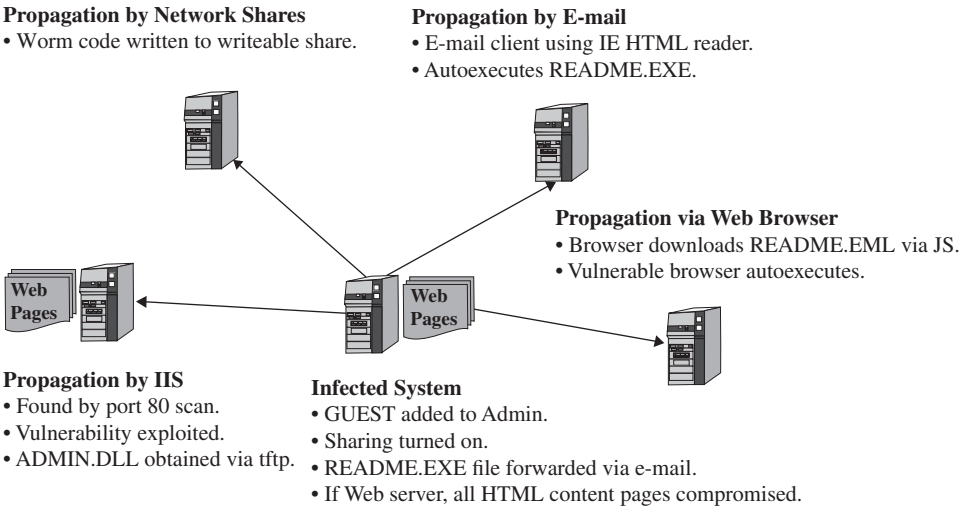


EXHIBIT 18.7 NIMDA Propagation Vectors

once asleep, the worm stayed asleep, although it could be reawakened. Removing the worm program from RAM required only a reboot, and the patch from Microsoft would prevent further infection.

As an aside, although only IIS servers could be exploited, many other devices that listen on port 80 were also affected. Cisco 600 DSL routers and HP JetDirect devices, for example, listen on port 80 and would crash when they received the buffer overflow packet.

Three different variants of Code Red existed on the Internet, all acting as described. In August 2001, a couple of new variants appeared that were called *Code Red II*. Unlike Code Red, Code Red II did not deface Web pages nor did it launch a DDoS attack on any given site. Instead, this worm was destructive, installing backdoors on infected servers, changing many registry settings, installing a Trojan horse version of *explorer.exe* (Windows Explorer), and disabling the System File Checker (SFC) utility. The worm also spread quickly, employing up to 300 threads at a time looking for other systems to infect.

18.8.2 NIMDA. The next evolution appeared in September 2001 and was called *NIMDA*. NIMDA (*admin* backward) was unique because it exploited multiple vulnerabilities in Microsoft code, namely IIS, Internet Explorer (IE), and the Message Application Program Interface (MAPI). As a result, NIMDA had four distinct propagation vectors (see Exhibit 18.7):

1. **IIS.** When a Web server was found, the attacker attempted to exploit various IIS vulnerabilities, including IIS sadmind, a Code Red II root.exe or other backdoor program, or IIS Directory Traversal. If successful, the attacker used the Trivial File Transfer Protocol (tftp) from cmd.exe to send the worm code (admin.dll) to the victim.
2. **Web browser.** The worm on an infected server created a copy of itself in a file called *readme.eml*. The worm also altered every Web-content file at the infected site with a small JavaScript code that pointed to this file. When a user browsed to

the infected Web server, the infected page's JS code was activated, and *readme.eml* was downloaded. Vulnerable versions of Internet Explorer would auto-execute the file while most other browsers would not.

3. **Email.** NIMDA sent itself to all of the email addresses found in the InBox and Address Book of an infected server in a MIME-encoded, 56KB attached file named *readme.exe*. The file had an "audio/x-wav" section that contained the worm. Email clients using IE 5.1 or earlier to display HTML would automatically execute the attachment if the message was opened or previewed.
4. **Network shares.** When on an infected system, the worm copied itself to all local directories on the victim host and to all open, writeable network shares. The worm also set up shares on the victim host.

In addition, the GUEST account on the infected system was activated and made a member of Administrator group. (Interestingly, after a summer of increasing DDoS worms, almost all such activity ceased after 9/11 for many months.)

Over the years, DDoS attacks have continued, but there is a dual purpose. Some DDoS attacks are intended to knock a server or network off the Net the old-fashioned way, namely by sucking up all of the bandwidth or other resource. Other attacks, however, use the target site's defenses against itself; by performing reconnaissance to understand how a site will respond to various stimuli, the attacker can actually send a crafted set of packets to the target site that will cause the target to self-limit bandwidth and/or acceptable source addresses.

18.9 DEFENSES AGAINST DISTRIBUTED DENIAL-OF-SERVICE ATTACKS.

As with DoS attacks, a site cannot in isolation defend itself from DDoS attacks. Members of the Internet community must work together to protect every site against becoming the source of attacks or forwarding the attacks. This section discusses some ways to help prevent the spread of DDoS attacks by limiting the distribution of the tools and by limiting the propagation of the offending attack packets.

Although not discussed in detail here, another point needs to be made about DDoS attack responses. As discussed in Chapter 53 in this *Handbook*, victims of such an attack should maintain detailed logs of all actions they take and events they detect. These logs may prove invaluable in understanding the attack, in preventing other attacks at the initial target and others, and in aiding law enforcement efforts to track down the perpetrators.

18.9.1 User and System Administrator Actions.

These steps should be taken to minimize the potential that an individual system will be compromised and attacked or used as a stepping-stone to attack others:

1. Keep abreast of the security vulnerabilities for all of the site's hardware, operating systems, and application and other software. This sounds like a Herculean task, but it is essential to safeguarding the network. Apply patches and updates as soon as possible. Standardize on certain hardware, operating systems, and software where feasible to help manage the problem.
2. Use host-based firewall software on workstations to detect an attack.
3. Monitor systems frequently to test for known operating system vulnerabilities. Regularly check to see what TCP/UDP ports are in use using the *netstat -a*

command; every open port should be associated with a known application. Turn off all unused applications.

4. Regularly monitor system logs and look for suspicious activity.
5. Use available tools to periodically audit systems, particularly servers, to ensure that there have been no unauthorized/unknown changes to the file system, registry, user account database, and so on.
6. Every client system must use anti-malware software that updates itself continuously to keep abreast of the constantly changing threat landscape. Effective anti-malware tools include integration with browsers for automatically blocking access to Websites that are known to harbor dangerous code. Such tools also integrate into email clients to block automatic opening of attachments. Users should not be permitted to disable the anti-malware software.
7. Do not download software from unknown, untrusted sites. If possible, know the author of the code. Even better, download source code, review it, and compile it on a trustworthy system rather than downloading binaries or executables.
8. Keep up with and follow recommendations from NIST, US-CERT, anti-virus research labs, hardware and software vendors, and other sources of best practices.

18.9.2 Local Network Actions. Even if users lock down their systems so that no vulnerability has gone unpatched and no exposure unprotected, the local network itself still can be at risk. Local network managers and network administrators can take several steps to protect all of their own users as well as the rest of the Internet community:

1. Every network connected to the Internet should perform egress address filtering at the router. *Egress filtering* means that the router should examine the Source Address field of every outgoing IP packet sent to the Internet to be sure that the NET_ID matches the NET_ID of the network. Historically, firewalls have been used to protect networks from attacks from the outside world. But those attacks come from somewhere, so sites should also use the firewall to *protect* the outside world.
2. Networks should block incoming packets addressed to the broadcast address (the all-ones HOST_ID). There is no legitimate reason that an external network device should be sending a broadcast message to every host on a network.
3. To prevent a site from being used as a broadcast amplification point, turn off the Directed Broadcast capability at the router unless it is absolutely essential. If it is essential, reexamine the network to see if there is not a better way or if the broadcast scope can be minimized. Even where Directed Broadcasts are useful, typically they are needed only *within* the enterprise and are not required for hosts on the outside.
4. RFC 1918 defines three blocks within the IP address space that are reserved for private IP networks; these addresses are not to be routed on the Internet.

IP Address Range	Network ID/ Subnet Mask	Number of Equivalent Classful IP Networks
10.0.0.0–10.255.255.255	10/8 prefix	1 Class A network
172.16.0.0–172.31.255.255	172.16/12 prefix	16 Class B networks
192.168.0.0–192.168.255.255	192.168/16 prefix	256 Class C networks

In addition, there are a number of reserved IP addresses per RFC 3330 that are never assigned to public networks or hosts, including:

0.0.0.0/32	Historical broadcast address
127.0.0.0/8	Loopback network identifier
169.254.0.0/16	Link-local networks
192.0.2.0/24	TEST-NET
224.0.0.0/4	Multicast address range
240.0.0.0/5	Reserved for future use
255.255.255.255/32	Broadcast

Attackers commonly use IP address spoofing, generally by using one of the RFC 1918 private addresses or one of the other reserved addresses. Firewalls should immediately discard *any* packet that contains any RFC 1918 or reserved IP address in the Source Address or Destination Address field; such packets should never be sent to the Internet.

5. Block all unused application ports at the firewall, particularly such ports as IRC (6665–6669/tcp) and those known to be associated with DDoS software.
6. Use stateful firewalls that can better examine a packet within the context of the entire packet exchange. Stateless firewalls only look at packets according to a simple set of rules but are not aware of the overall packet traffic on the network (e.g., a stateless packet filter might forward an ICMP *echo* Reply packet to be sent out because the rules allow it; a stateful packet filter will forward the packet only if the rule allows it *and* there was a corresponding ICMP *echo* Request).
7. Use intrusion detection and intrusion prevention to protect the network. For example, personal firewall software can be installed on every workstation to help detect an attack on individual systems; this strategy is particularly useful at sites that have a large number of systems *in front* of a firewall (e.g., colleges). It is no coincidence that so many daemons reside on college and university computers that have been *0wned* (i.e., taken over by hackers).
8. Regularly monitor network activity so that aberrations in traffic flow can be detected quickly.
9. Educate users about events to watch for on their systems and how to report any irregularity that might indicate that someone or something has tampered with their system. Educate the help desk and technical support to assist those users who make such reports. Have an intelligence-gathering system within the organization so that such reports can be coordinated centrally to spot trends and to devise responses.
10. Follow NIST, US-CERT, and other best practices procedures.

18.9.3 Internet Service Provider Actions. ISPs offer the last hope in defeating the spread of a DDoS attack. Although the ISP cannot take responsibility for locking down every customer's host systems, ISPs have—and should accept—the responsibility to ensure that their network does not carry packets that contain obviously "bad" packets. Some of the steps that ISPs can take include:

1. As mentioned, attackers commonly employ IP address spoofing using an RFC 1918 private address or other reserved address. Amazingly, many ISPs will route

these packets. Indeed, there is no entry in their routing table telling them where to send the packets; they merely forward them to a default upstream ISP. *Any* packet that contains any RFC 1918 or reserved IP address in the IP Source Address or Destination Address field should be discarded immediately.

2. Perform ingress (and egress) address filtering. *Ingress filtering* means that ISPs should examine every incoming packet to their network from a customer's site and examine the IP Source Address field to be sure that the NET_ID matches the NET_ID assigned to that customer. Doing this will require additional configuration at the router and may even result in slight performance degradation, but the trade-off is certainly well worth the effort. The ISPs also should perform egress filtering to check their outbound packets to upstream and peer ISPs.
3. Disable IP directed broadcasts.
4. Pay careful attention to high-profile systems (servers) and customers.
5. Educate customers about security and work with them to help protect themselves.

Most of the ISP community takes at least some of these steps. Users should insist that their ISPs provide at least these protections and should not do business with those that do not. RFC 3013 and the North American Network Operators' Group (NANOG) are good sources of information for ISPs.

18.9.4 Exploited Software Defensive Actions. There are a number of defensive steps that can be taken to avoid or mitigate problems due to Code Red/NIMDA-like attacks that exploit software. Several of these recommendations are controversial due to their implicit boycott of products from a specific vendor.

- If using IIS, consider using alternate Web server software. If use of IIS is essential, keep IIS and the operating system up to the latest patch revision. Microsoft's IIS Cumulative Patch, for example, does *not* clean a system of many of the backdoors used for exploits.
- If using Internet Explorer, consider using alternate browser software. If you must use IE, secure it against MIME auto-execution. Note that as late as September 2012, unpatched vulnerabilities were being reported by US-CERT in IE6, IE7, IE8, and IE9, and some authorities (e.g., the government of Germany) were recommending against the use of any version of IE.
- Disable all unused accounts on servers and other systems. In particular, enable the Guest account or anonymous access only if absolutely necessary.
- Disable JavaScript, Java, and ActiveX on browsers unless absolutely necessary.
- Do not execute or open *any* email attachment unless expected, known, and verified.
- Use the most up-to-date antivirus signature files.
- Unbind file and print sharing from TCP/IP. In some cases, this will require installing NetBEUI for file and print sharing.

18.9.5 Other Tools under Development or Consideration. Responses to DDoS attacks are not limited to the defensive steps just listed. Indeed, proactive responses to the prevention and detection of DDoS attacks are an active area of research. These proposals are merely samples of some ways for dealing with DDoS attacks; the first adds new hardware to the Internet, the second requires changing Web server and

client software, and the latter two require incrementally changing software in all of the Internet's routers and hosts, respectively. Upgrading Web browsers is probably the most practical strategy even though there are millions of copies in distribution; the vast majority come from just a handful of vendors, and users tend to upgrade frequently anyway.

18.9.5.1 Distributed Traffic Monitor. One method that has be proposed is to examine the network at the ISP level and build a type of intelligent, distributed network traffic monitor; in some sense, this would be like an IDS for the Internet. ISPs, peering points, and/or major host servers would have traffic monitor hardware using IP and the Internet for communications, much like today's routing protocols. Each node would examine packets and their contents, doing a statistical analysis of traffic to learn the normal patterns. These devices would have enough intelligence to be able to detect changes in traffic level and determine whether those changes reflected a normal condition or not. As an example, suppose that such hardware at Amazon.com were to identify a DoS attack launched from an ISP in Gondwanaland; the traffic-monitoring network would shut off traffic to Amazon coming from that ISP as close to the ISP as possible. In this way, the distributed network of monitors could shut traffic off at the source.

The hardware would need to be informed about traffic-level changes due to normal events, such as a new Super Bowl commercial being posted on YouTube or a new fashion show at Victoria Secret's Website. The hardware also would need to prevent the attacker community from operating under the cover of these normal events.

18.9.5.2 Client Puzzle Protocol. RSA Laboratories has proposed cryptographic methods as a potential defense to DDoS attacks against Web servers. This approach would use a *client puzzle* protocol designed to allow servers to accept connection requests from legitimate clients and block those from attackers. A client puzzle is a cryptographic problem that is generated in such a way as to be dependent on time and information unique to the server and client request.

Under normal conditions, a server accepts any connection request from any client. If an attack is detected, the server selectively accepts connection requests by responding to each request with a puzzle. The server allocates the resources necessary to support a connection only to those clients that respond correctly to the puzzle within some regular TCP time-out period. A bona fide client will experience only a modest delay getting a connection during an attack, while the attacker will expend an incredible amount of processing power to continue sending the number of requests necessary for a noticeable interruption in service at the target site, quickly rendering the attack ineffective (in effect, a reverse DoS). This scheme might be effective against a DDoS attack from a relatively small number of hosts each sending a high volume of packets but might have limited effectiveness against a low-volume attack from a large number of systems or from a botnet.

18.9.5.3 IP Traceback. A third mechanism that was the subject of considerable research during the early to mid-2000s was *IP Traceback*. The problem with DoS/DDoS attacks is that packets come from a large number of sources, and IP address spoofing masks those sources. Traceback marking, in concept, is a relatively straightforward idea. Every packet on the Internet goes through some number of ISP routers. Processing power, memory, and storage are available for routers to mark packets with partial path information as they arrive. Since DoS/DDoS attacks generally comprise a large number

of packets, the traceback mechanism does not need to mark every packet, but only a sample size that is statistically likely to include attack packets (e.g., 1 packet out of every 20,000, or 0.005% of the IP traffic). This feature would allow the victim to locate the approximate source of the attack without the aid of outside agencies and even after the attack had ended. Another traceback proposal would define an ICMP Traceback message that would be sent to the victim site containing partial route information about the sampled packet. There are many issues related to traceback that need to be resolved, such as the minimum number of marked packets required to reconstruct the path back to the attacker, the actual processing overhead, and the ability to perform traceback while an attack is under way. In addition, any traceback solution will require a change to tens of thousands of routers in the Internet; how effective can traceback be during a period of gradual implementation? The upside, of course, is that the solution is backward compatible and results in no negative effects on users.

18.9.5.4 Host Identity Payload. A fourth proposal was to modify IP to be less prone to address spoofing by making the protocol less dependent on the address field for anything more than routing. The Host Identity Payload (HIP), for example, defines a protocol for the exchange of a cryptographic *Host Identity* between two communicating systems. This feature relegates the IP address for use solely as a mechanism for packet forwarding rather than as an identifier of the sender. Sender identification, instead, is accomplished by the Host Identity value, and all of the higher layer protocols are bound to the Host Identity. HIP is not yet widely deployed but is available in some TCP/IP implementations.

18.10 MANAGEMENT ISSUES. One of the greatest shortcomings in many organizations is that the highest levels of management do not truly understand the critical role that computers, networks, information, and the Internet play in the life of the organization. It is difficult to explain that there is an intruder community that is actively working on new tools all the time; and history has shown that as the tools mature and become more sophisticated, the technical knowledge required of the potential attacker goes down and the number of attacks overall goes up. Too many companies insist that "no one would bother us" without realizing that *any site* can become a target just by being there.

DoS attacks come in a variety of forms and aim at a variety of services, causing increased complexity and difficulty for system defense. DoS attacks should be taken seriously because of the potential threat they present, and attempts should be made to educate operational staff before such attacks occur, to document DoS attacks if they do occur, and to review the documentation and actions taken after the incident is over. Discussion of what steps were taken, what actions went into effect, and what the overall result was will help in determining whether the procedures carried out and techniques utilized were those best suited to the situation. A frank review and discussion will help achieve the best, most rapid, and most effective deployment of resources.

If anything proves the intertwined nature of the Internet, it is the defense against DDoS attacks. DDoS attacks require the subversion and coordination of hundreds or thousands of computers to attack a few victims. Defense against DDoS attacks requires the cooperation of thousands of ISPs and customer networks. Fighting DDoS requires continued diligence in locking down all of the hosts connected to the Internet as well as fundamental changes in the nature of TCP/IP connection protocols. In addition, many of the same techniques used to insert viruses and worms that lead to DoS and DDoS attacks are being used to insert Advanced Persistent Threat (APT)–class malware.

As with most security concerns on the Internet, a large part of the solution lies in user education. Users must be educated as to what current threats are currently on the Internet and affecting relevant applications, what to do if something suspicious occurs, and how to respond to attackers that are increasingly creative. The world of everyone having their own computers, mobile devices with an incredible amount of intelligence, and a bring-your-own-device (BYOD) corporate atmosphere requires that every user in the enterprise is a potential vulnerable point and needs to be properly educated.

18.11 FURTHER READING

Anonymous. *Maximum Security*, 4th ed. Indianapolis: SAMS, 2003.

CERT Coordination Center. *Results of the Distributed-Systems Intruder Tools Workshop*. CERT Website, December 7, 1999. www.cert.org/reports/dsit_workshop.pdf

Denning, D. E. *Information Warfare and Security*. Reading, MA: Addison-Wesley, 1999.

Ferguson, P., and D. Senie. "Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing" (RFC 2827/BCP 38). The Internet Society, May 2000. www.ietf.org/rfc/rfc2827.txt

Gao, Z., and N. Ansari. " Tracing Cyber Attacks from the Practical Perspective." *IEEE Communications Magazine* 43, no. 5 (May 2005): 123–131.

Gibson, S. "The Strange Tale of the Denial of Service Attacks Against GRC.COM." Computer Crime Research Center Website, June 2, 2001. www.crime-research.org/library/grcdos.pdf

McClure, S., J. Scambray, and G. Kurtz. *Hacking Exposed, Network Security Secrets & Solutions*, 6th ed. Berkeley, CA: Osborne/McGraw-Hill, 2009.

Mirkovic, Jelena, Sven Dietrich, David Dittrich, and Peter Reiher. *Internet Denial of Service: Attack and Defense Mechanisms*. Upper Saddle River, NJ: Prentice-Hall, 2005.

Raghavan, S. V., and E. Dawson, eds. *An Investigation into the Detection and Mitigation of Denial of Service (DoS) Attacks: Critical Information Infrastructure Protection.* Springer, 2011.

Rekhter, Y., B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. "Address Allocation for Private Internets" (RFC 1918/BCP 5). The Internet Society, February 1996. www.ietf.org/rfc/rfc1918.txt

Skoudis, E. *Counter Hack Reloaded: A Step-by-Step Guide to Computer attacks and Effective Defenses*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 2006.

Spafford, E. H. " The Internet Worm: An Analysis." *Computer Communication Review* 19, no. 1 (January 1989): 17–57.

Most of the best current references are on the Web. Some of the sites that readers should monitor for DoS/DDoS information are:

ATLAS Summary Report, Global Denial of Service: http://atlas.arbor.net/summary/dos

CERT/CC: www.cert.org

Dave Dittrich's "Distributed Denial of Service DDoS Attacks/tools" Web page: http://staff.washington.edu/dittrich/misc/ddos/

Denialinfo.com's "Denial of Service (DoS) Attack Resources" Web page: www.denialinfo.com/dos.html

DShield: www.dshield.org

SANS Internet Storm Center: http://isc.sans.org

US-CERT: www.us-cert.gov

18.12 NOTES

1. DHS National Cyber Security Division, US-CERT, "National Vulnerability Database," Website, 2013, http://web.nvd.nist.gov/view/vuln/search?execution=e2s1
2. Formerly Internet Information Server.
3. 10 KB × 8 bits/byte × 50 hosts = 4 Mb.
4. Michael Calce and Craig Silverman, *Mafiaboy: A Portrait of the Hacker as a Young Man*, Lyons Press, 2011.
5. Brian Bloom, "The New DDoS: Silent, Organized, and Profitable," *PCWorld*, May 29, 2012, www.pcworld.com/article/256431/the_new_ddos_silent_organized_and_profitable.html

CHAPTER 19

SOCIAL-ENGINEERING AND LOW-TECH ATTACKS

Karthik Raman, Susan Baumes, Kevin Beets, and Carl Ness

19.1 INTRODUCTION 19·1

19.2 BACKGROUND AND HISTORY 19·2

19.3 SOCIAL-ENGINEERING METHODS 19·4

19.3.1 Impersonation 19·4
19.3.2 Seduction 19·5
19.3.3 Intimidation 19·6
19.3.4 Low-Tech Attacks 19·6
19.3.5 Network and Voice Methods 19·8
19.3.6 Reverse Social Engineering 19·11

19.4 THE PSYCHOLOGY OF SOCIAL ENGINEERING 19·11

19.4.1 Introduction 19·11
19.4.2 Psychology 19·12
19.4.3 Social Psychology 19·13
19.4.4 The Social Engineer Profile 19·14

19.5 DANGERS OF SOCIAL ENGINEERING AND ITS IMPACT ON BUSINESSES 19·14

19.5.1 Consequences 19·14
19.5.2 Case Studies—Examples from Business 19·15
19.5.3 Success Rate 19·16
19.5.4 Small Business verses Large Organizations 19·16
19.5.5 Trends 19·17

19.6 DETECTION 19·17

19.6.1 People 19·17
19.6.2 Audit Controls and Event Logging 19·18
19.6.3 Technology for Detection 19·18

19.7 RESPONSE 19·19

19.8 DEFENSE AND MITIGATION 19·19

19.8.1 Training and Awareness 19·20
19.8.2 Technology for Prevention 19·20
19.8.3 Physical Security 19·21

19.9 CONCLUDING REMARKS 19·21

19.10 NOTES 19·22

19.1 INTRODUCTION. According to Greek mythology, the Greeks defeated the Trojans in the Trojan War with the help of a wooden statue. After fighting a decade-long war in vain, the Greeks withdrew from their stronghold on the beach. Outside the gates of Troy, they left a giant wooden horse. The statue confused the Trojan soldiers, but it was brought within the fortified walls of Troy. Inside the statue hid several Greek soldiers. When darkness fell, these soldiers emerged from the statue and opened the

gates of Troy. The Greek army entered Troy and took the soldiers and citizens of Troy by surprise. After the deception, Greece won the war quickly.

The Trojan horse built by the Greeks was effective because it used deception to achieve the desired result: penetrating the enemy's established defenses. The Trojan horse accomplished what those in information security label *social engineering*.

Social engineering may be defined as obtaining information or resources from victims using coercion or deceit.

> Social engineering refers to lying, cheating, tricking, seducing, extorting, intimidating, and even threatening employees into revealing confidential information that can then be used to break into systems. Social engineering is based on deception and on violation of social norms of fairness and honesty.[1]

During a social-engineering attack, attackers do not depend on subverting technology, for example, scanning networks, cracking passwords using brute force, or exploiting software vulnerabilities. Rather, social engineers operate in the social world by manipulating the trust or gullibility of human beings. Thus, social engineering relies on manipulating human nature to extract information or access by tricking a victim. Related to social engineering are low-tech attacks, and the two often go hand-in-hand. Low-tech attacks are similar to social-engineering attacks in that they do not utilize technology. They are physical attacks performed against companies or individuals' property.

Not all social-engineering and low-tech attacks will give attackers all the information they are seeking at once. Social engineers will collect small pieces of information that seem innocuous to the individuals that divulge them. Social engineers may gather these snippets of information in a seemingly random order but then assemble them into intelligence that is used to launch larger attacks that can be devastating to an organization's information security, resources, finances, reputation, or competitive advantage. Indeed, the purpose of a social-engineering attack can be as varied as the attack method employed. The result, however, is generally the same: a loss of intellectual property, money, business advantage, credibility, or all of the above.

Social-engineering attacks have been and will remain successful due to the weakest security link in an organization: the people. They are important because they exploit human nature, which is immutable and therefore perpetually vulnerable.

This chapter presents the history of social-engineering and low-tech attacks, its methods, the social science behind it, and its business impact. In addition, it covers detection and mitigation policies for managers and information security officers to defend against and mitigate social-engineering and low-tech attacks.

The purpose of a social-engineering attack can be as diverse as the attack method. Nevertheless, the result to the victims is generally the same, loss of intellectual property, money or business, credibility, or all of the previous. This chapter presents the history of social-engineering and low-tech attacks, its methods, the social science behind it, and its business impact. In addition, it covers detection and mitigation policies for managers and information security officers to defend against social-engineering and low-tech attacks.

19.2 BACKGROUND AND HISTORY. Social engineering is not a new tactic, nor is it an invention of modern-day hackers. The term has its foundations in political history, where a person or group manipulates a group of people, large or small, in an attempt to persuade or manipulate social attitudes or beliefs. Often, governments or political parties engage in this practice. The modern term's origins date back before World War II began, and the Nazis are thought to have actually coined the term.[2] To this day,

the term carries a negative connotation because of its roots in history, especially during Nazi-controlled Germany. Some researchers also consider the realm of social engineering to cover everything from advertising and modern media to political action groups.

Social engineering today is better known for its use as a security-penetration technique.

Deception has been integral to espionage through the centuries. Sun Tzu wrote in the fifth century BCE,

- "All warfare is deception."
- "Having doomed spies, doing certain things openly for purposes of deception, and allowing our spies to know of them and report them to the enemy."[3]

Deception—specifically, feeding false information to the enemy's spies—was an integral part of the Second Punic War between Rome and Carthage in the third century BCE.[4] During the Second World War, the Allies used "Operation Bodyguard" to deceive the Axis powers into believing that the D-Day invasion would occur at a different time and place from its true schedule and target.[5] Deception was integral to the spy vs. counter-spy competition between the Communist block and the West during the Cold War.[6]

A well-known example of social engineering is the escapades of Frank W. Abagnale, subject of the fictionalized movie *Catch Me If You Can*.[7] Abagnale was able to successfully impersonate authority figures, including a physician, pilot, attorney, and teacher, even though he was a teenager. He also used social-engineering techniques to persuade and manipulate innocent and good-natured individuals to help him carry out many of his frauds. Many of Abagnale's techniques were highly successful and well engineered. Today, he helps organizations recognize and defend against such attacks through his speaking engagements and consulting business.[8]

One of the best-known social engineers is Kevin Mitnick. Though Mitnick is now a computer security consultant, lecturer, and author, his present career was preceded by a number of years spent as a computer hacker and social engineer. Long reformed, Mitnick has written several books discussing his observations and techniques as a computer hacker. Mitnick maintains that social engineering is the most powerful tool in the hacker's toolbox.

Ever since social engineering has become a popular, and more importantly, successful, technique, its frequency of use has increased. One of the most visible reminders of this is the large number of phishing and pharming attacks, discussed in detail in Chapter 20 in this *Handbook*.

Social engineering has even become a spectator sport at some criminal-hacker conventions. At the 2012 DEF CON in Las Vegas, Shane MacDougall won the Social-Engineering Capture-the-Flag Contest by tricking a Walmart executive in a Canadian store into revealing every single piece of confidential information listed in the contest objectives.

> Darnell asked the manager about all of his store's physical logistics: its janitorial contractor, cafeteria food-services provider, employee pay cycle, and staff shift schedules. He learned what time the managers take their breaks and where they usually go for lunch.
>
> Keeping up a steady patter about the new project and life in Bentonville, Darnell got the manager to give up some key details about the type of PC he used. Darnell quickly found out the make and version numbers of the computer's operating system, Web browser, and antivirus software.[9]

Social engineering, when used by itself as a standalone attack tool, is very effective. But this tool can be used as part of a larger attack, or even as part of a technical attack. Social engineering is commonly used at the beginning of a larger, more substantial attack. It is important to remember that some attacks have taken place over weeks, months, or years. This is not always the case, and just because a specific social-engineering attempt was averted, the larger attack may not be completely mitigated. An example of this may be that an initial attack was unsuccessful because the attacker did not have sufficient or accurate information to carry out the attack. The attacker may revert once again to social engineering and low-tech attack strategies to gather information for a different, possibly successful, technical attack.

Even though an organization may have the best firewall, intrusion detection system, and risk management tools in the world, an attacker may use social engineering to circumvent these technical defenses. The attacker will not be stopped by the best technical defenses if the attacker is able to extract a valid user name and password from an unsuspecting employee. Once the attacker has this information, it may be enough information to carry out a massive attack on an organization's information systems—possibly undetected.

Computer criminals have been able to use pure social engineering and low-tech attack techniques, without relying on any serious technology, to cause large-scale damage to an organization as well. An action as simple as taking discarded, potentially damaging, information out of a dumpster and sending it to local media can cause substantial damage to an organization's image. This act might be part of a larger campaign or operation by a group of individuals seeking political or social action against an organization. An example of this was recently chronicled by Home Box Office (HBO) involving a group of citizens who routinely gathered the trash from county election centers and city halls in an effort to collect evidence that they could present to the media to prove voter fraud. Documents that were improperly disposed of have had extremely negative consequences for election officials.[10]

19.3 SOCIAL-ENGINEERING METHODS. Social-engineering attacks can take many different forms and expert social engineers are capable of changing their methods of attack very quickly in order to succeed. The underlying principle of most attacks is *pretexting*, which is defined as "the collection of information … under false pretenses."[11] Two distinct underlying methodologies are used during social-engineering attacks, impersonation and seduction. The basis for most attacks is either of these two methods. Certain low-tech attacks, which do not involve any human contact, are exceptions.

Targets of social-engineering attacks also vary widely. Depending on how complex the attack is or how much knowledge the attacker has will help determine the target for the attack. However, in many instances, a social-engineering attack is an attack of opportunity and the victim will be randomly chosen. In certain well-planned attacks, a particular target may be identified as the victim.

19.3.1 Impersonation. Impersonation is defined as pretending to be someone else, and it is one of the most popular methods that social engineers employ. Social engineers may use cross-functional impersonation attacks to target an employee at any level within the targeted organization. They may pretend to be a real, named individual, or they may pretend to have a particular role or authority.

Corporate executives are common targets of social engineering.[12] Helpdesk employees and systems administrators are also common targets of impersonation. Most

organizations have helpdesks for assisting employees with issues related to information technology (IT). Employees generally will follow the instructions from helpdesk personnel simply because they are perceived as technologically well-informed and trustworthy. Social engineers understand this trust and will exploit it to steal information. The attacker will impersonate helpdesk personnel and abuse this blind trust to gather whatever information the attacker needs.

Helpdesk personnel can also be the victims of social-engineering attacks with the social engineer impersonating a user in need of technical assistance. An attack against the AOL helpdesk is a historical example of a successful social-engineering attack against helpdesk personnel. An attacker posing as a customer was able to infiltrate the AOL environment by conning an AOL helpdesk staff member.[13]

There have also been instances where social engineers impersonate corporate officers or managers. In one case, the payroll giant ADP released personnel and brokerage account information on hundreds of thousands of customers. The attacker impersonated a corporate board member, requested and received the information.[14] This case highlights the need for controls and education for all levels of employees. The victim in this case might have been concerned about not pleasing a powerful member of the organization and may have feared retribution.

It is important to note that attackers will also impersonate regular employees of an organization by dressing, speaking, and blending into the organization's environment. In doing so, the attacker may gain physical access to a facility by piggybacking or tailgating. In one case reported by a student in a 1993 security course, a social engineer collaborating with a corrupt manager was able to commandeer an unused desk in a large company for several *months* before a security guard asked for his company identification. The criminal disappeared immediately, but not before ingratiating himself to the employees as a new employee and even being invited out to social events by the unsuspecting victims.[15]

For a social engineer, there are no boundaries for impersonation; they may try to gain information via physical access to an organization using any number of ruses including impersonation of:

- Temporary employees (e.g., contractors or auditors)
- Utility or telecommunications company employees
- Emergency personnel
- Janitorial or maintenance employees
- New employees
- Delivery personnel

Physical access can greatly increase the success rate of an attack depending on whether the organization has proper controls in place. Highly motivated social engineers may go so far as to gain employment at the company, or at a company that is a client of the company, in order to have easier access to the victim.

19.3.2 Seduction. The word seduce is defined as "to lead away from duty, accepted principles, or proper conduct".[16] In general, a social-engineering attack using seduction will take longer to complete then an impersonation attack. The attacker, using seduction, will identify a target and will form a bond with that individual, through social settings, online, or through another mechanism. It is during this relationship that the victim's information is divulged to the attacker.

For example, a social engineer who wishes to gain access to a building may befriend a security guard of that organization. After some time has passed and the relationship has progressed, the attacker may request a tour of the facility. The security guard, wanting to please the new friend, may allow a tour. The social engineer, once inside, can plant clandestine listening devices, look for usernames or passwords, and read documents left in the open.

19.3.3 Intimidation. Social engineers can use fear to achieve their ends. For example, in 2012, scammers spread alarming emails about FB I prosecutions to innocent victims.

> A new extortion technique is being deployed by cybercriminals using the Citadel malware platform to deliver Reveton ransomware. The latest version of the ransomware uses the name of the Internet Crime Complaint Center to frighten victims into sending money to the perpetrators. In addition to instilling a fear of prosecution, this version of the malware also claims that the user's computer activity is being recorded using audio, video, and other devices.
>
> As described in prior alerts on this malware, it lures the victim to a drive-by download Website, at which time the ransomware is installed on the user's computer. Once installed, the computer freezes and a screen is displayed warning the user they have violated United States Federal Law. The message further declares that a law enforcement agency has determined that a computer using the victim's IP address has accessed child pornography and other illegal content.
>
> To unlock the computer, the user is instructed to pay a fine using prepaid money card services. The geographic location of the user's PC determines what payment services are offered. In addition to the ransomware, the Citadel malware continues to operate on the compromised computer and can be used to commit online banking and credit card fraud.[17]

Another social-engineering attack using intimidation is the phony office-supply scam.[18] A low-level employee accepts an offer of free printer toner, glass cleaner, or paper and receives a box of materials—followed by an invoice. Attempts to return the unwanted, no-longer-free materials are refused and the criminals use increasingly strong threats to intimidate the employee into passing the invoice along for payment. Employees may be told that their employer will be sued or that a complaint about their behavior will be sent to upper echelons of the organization. A variation involves sending a gift to the victim in the hope that accepting it can be used to embarrass the recipient into complying with the scammers' demands. Such scams succeed often enough to generate millions of dollars of revenue for the criminals.[19]

19.3.4 Low-Tech Attacks. Low-tech attacks are invaluable to an attacker as part of reconnaissance, information theft, and surveying for easy targets. On occasion, these methods may even reward the attacker with a very large amount of useful information in a very short amount of time. Low-tech attack methods may seem simple or improbable, but the methods described can easily be overlooked by security managers. They are not urban legends—they have been used in the past and continue to be utilized today.

19.3.4.1 Dumpster® Diving. In the context of social engineering, *Dumpster® diving* is the social engineer's act of searching through an organization's garbage in an attempt to find documents, hardware, software, or anything that could be of value to meet the goals of the attacker. Even with the widening use of sensitive document destruction, Dumpster® diving is a popular social-engineering technique because it is easy and often successful. Oracle hired detectives to purchase Microsoft's trash during

Microsoft's antitrust trial. (The detectives were unsuccessful.)[20] Social engineers do not need to deceive anyone to perform the attack. In many cases, the materials disposed may sit in open containers for weeks. Dumpster® diving is most often carried out at night when no one is around, as there is less risk of being caught. So as to not draw attention to themselves, Dumpster® divers have been known to dress in dark clothing or even use janitorial uniforms.

All organizations must understand the legal ramifications of Dumpster® diving. Local and state laws may vary widely and be murky; however, no organization or individual should have *any* realistic expectation of privacy relating to materials in refuse containers. It may even be legal for an attacker to remove and take ownership of anything left in a garbage receptacle if it is placed outside the limits of private property. Confidential information must be destroyed, not simply discarded.

An unfortunate example of making confidential data accessible in garbage came to light in November 2012, when strips of confetti thrown into the crowd at the Macy's Thanksgiving parade were discovered to have intact, readable confidential data:

> Among the easily identifiable records from the Nassau County Police Department were what appears to be details of Mitt Romney's motorcade route to and from the final presidential debate at Hofstra University.
>
> Confetti collected by spectators near 65th Street and Central Park West also contained arrest records, incident reports, and personal information, and identified undercover officers, WPIX-TV says.
>
> "There are phone numbers, addresses, more Social Security numbers, license plate numbers," said Ethan Finkelstein, 18, of Manhattan, who gathered up some of the confetti with friends. "And then we find all these incident reports from police."[21]

The papers had been shredded the wrong way in a single-cut shredder, leaving horizontal strips of readable data for anyone to see.

19.3.4.2 Theft. The age-old crime of theft is another popular social-engineering technique. Social engineers may pick up anything they can get their hands on, literally, and leverage the information obtained to carry out other attacks. Targets of theft include, but are not limited to, printed materials, CD-ROMs, USB flash drives, backup media, tablets, smart devices, and laptops. Thieves may obtain objects on or off company premises. While on company premises, they may look for objects they can grab and quickly conceal. An attacker may bring in empty laptop cases, backpacks, or even large purses to aid in their efforts. Most employees are more likely aware of theft techniques outside of the organization, such as taxicabs, airports, and other public places. Employees may not realize that a social engineer or other criminal may also be an insider.

19.3.4.3 Leveraging Social Settings. Social engineers may use social settings to gain information because people relax in social settings and may believe that information security practices are for the workplace only. A social engineer may use a social setting, such as a bar, to take advantage of drinking employees to gain information. The attacker may actively engage a target or passively eavesdrop on a conversation. While this type of attack may seem far-fetched, there are many situations where the attacker may be in the right place at the right time to gain knowledge that can be later used as part of a larger attack. People in social settings are less likely to have their defenses up and security on their minds.

Restaurants, corporate functions, impromptu meetings outside of a company building, or loud phone conversations are all areas or situations where eavesdropping can occur. Many times, employees will work on commuter trains or conduct business in other public areas. A social engineer can exploit an organization's mobile workforce to gain information. Employees should be cognizant of who is around them while performing any work-related task.

19.3.4.4 Exploiting Curiosity or Naïveté. Social engineers may trick a victim into unknowingly aiding in an attack by piquing the user's curiosity. For example, an attacker may leave an intriguingly labeled CD-ROM in a break room, hoping that a victim is curious about the contents. When the victim places the CD-ROM in his or her computer, a malicious program, such as a virus, could automatically execute and spread. This technique has also been executed using USB drives, iPods, and corrupted music CDs.[22]

19.3.4.5 Data Mining and Data Grinding. Search engines can catalogue a surprising amount of information that is sensitive and confidential. Social engineers can create special searches, use search engine application programming interfaces (APIs), and use advanced search capabilities of many search engines to mine information about a company. Another attack vector is the caching feature of many search engines. A search engine may cache a Web page with sensitive information. If an organization requests the cache be removed, still there is a delay before the search engine cache is updated, during which time the organization's information is exposed. Social engineers could also exploit data about individuals or organizations gathered from social media. See Section 19.3.5.6, "Social Media," in this *Handbook* for details about the use of this method. Attackers could mine this data for victims' behavioral trends. They could garner more information about the victims than was readily presented by social media data.

Documents published by a company are another source of unintended information disclosure. In a technique known as data grinding, social engineers can use metadata-reader software to extract information such as the author's name, organization, computer name, network name, email address, user identification (ID), and comments from Microsoft Office documents.[23] This potentially damaging information is included in most document types.

19.3.4.6 Piggybacking or Tailgating. A common and very successful method of social engineering is *piggybacking* or *tailgating*. The method allows an attacker access to a facility without showing proper credentials. The victim in these cases is being polite and holding the door for the attacker who enters the facility using the credentials of the victim. Once inside a facility, an attacker is free to roam the building in search for information. If questioned upon entering, attackers may use the excuse of forgetting their credentials or being new employees. In general, if a piggybacking is going to be attempted, social engineers will dress and act according to other members of the organization so they can blend into the facility. If a facility is secured by multiple layers of physical security, attackers may try to use social engineering to breach the first few layers then proceed to use other attack methods.

19.3.5 Network and Voice Methods. The Internet has taken off, and so have social-engineering attacks. These attacks differ from traditional social-engineering attacks due to minimal or no human interaction. Still, the basis of the attacks still relies

on the trusting nature of humans. The annoying Nigerian 419 attack is a combination attack using both email and human interaction.[24]

The methods used for these attacks include *phishing*, *pharming*, *spim*, *malware*, and *vishing*. Phishing and Trojans are discussed in detail in Chapter 20 in this *Handbook*.

19.3.5.1 Phishing. Phishing is one of the most widely used and successful social-engineering attacks. It is defined as the "act of sending an email to a user falsely claiming to be an established legitimate enterprise in an attempt to scam the user into surrendering private information."[25] Phishing, as with all social-engineering attacks, relies on the trusting nature of people. Naïveté about using the Internet also plays a role in the success of this attack.

An interesting new wrinkle is that being *misrepresented as the origin* of phishing emails was the second most common type of attack reported by respondents (39 percent) to the Computer Security Institute's "2010/2011 Computer Crime and Security Survey."[26]

Newer, more targeted phishing attacks called *spear phishing* are also gaining in popularity. Spear phishing is used in a dangerous type of threats known as Advanced Persistent Threats (APTs). Spear phishing was used in a portion of a breach of the security firm RSA in 2011.[27]

19.3.5.2 Spim. Many people and businesses rely on the synchronous communication that instant messaging (IM) offers. Social engineers have noted the increase in the use of IM software and developed *spim*. Spim is "instant spam or IM spam."[28] A spim attack is very similar to a phishing attack except the vector is IM software instead of email. For example, an attacker will develop a fraudulent Website that resembles a legitimate one and send its link to many IM accounts. Victims will visit the Website and login, revealing their credentials to the attacker. The fraudulent site could contain malicious software that infects the victim's computer. Surprisingly, despite the increasing use of IM software, the growth of spim has been slow.[29]

19.3.5.3 Pharming. In *pharming* attacks, attackers attempt to make victims visit spoofed Websites to reveal sensitive personal information. Attackers achieve this by manipulating the victim's local or global DNS directory. Pharming attacks may fool users more easily than other low-tech attacks because the user may be given no indication that an attack is under way. Users may type in the URLs of their banking or credit card Websites into their browsers as usual and not normally notice that they are in fact visiting fraudulent sites. In 1997, Dan Parisi registered the domain whitehouse.com and snared people mistakenly looking for whitehouse.gov into viewing pornography—an early form of pharming.[30] In a 2005 pharming attack, users of the Internet mail service Hushmail were redirected to a fraudulent site where their information was harvested. Hushmail suffered negative publicity from this attack and was forced to update its users daily about its investigation into the attack.[31]

19.3.5.4 Malware. Malware are programs or files that are harmful or dangerous to the end user. Social engineers frequently use Trojans and viruses. Although the programs themselves do not make up a pure social-engineering attack, the basis of the attack still relies on manipulating a victim's trust. For example, a social engineer may send a victim a phishing email with a link to a malicious Website. If the user visits the malicious Website, a Trojan is installed on the victim's computer and the malicious program will begin gathering the victim's information. Another example involves a

USB drive loaded with auto-executing malware. A victim curious about the contents will plug the drive into a machine and execute the malicious code.[32] Any number of vectors can be used, including documents, email messages, Websites, CDs, and USB drives. The premise is the same for all types of these attacks: The unsuspecting users, curious about the content, will inadvertently install these dangerous programs. Malware infections were the most common type of attack reported by respondents (67 percent) to the Computer Security Institute's "2010/2011 Computer Crime and Security Survey," and bots or zombies were reported by 29 percent.[33]

Malware has recently begun to use a type of social-engineering attack whereby the trust placed in a valid certificate is exploited in conjunction with a long-known vulnerability called Data Link Library (DLL) preloading. In this attack, a package is prepared that includes a valid, trusted application that has a valid signature. In this package, however, a malicious DLL is included. Most users, who may know to look for a valid certificate, will indeed find one in the application and therefore trust the entire package. The victim launches the valid, trusted application, and due to the preloading vulnerability, the malicious DLL is loaded from the local directory before the valid DLL is loaded from its regular path.[34]

See Chapters 16, 17, and 18 in this *Handbook* for details of malware.

19.3.5.5 Vishing. Attackers who lure victims with the use of email and the telephone or just the telephone are performing a *vishing* attack.[35] Social engineers may send an email or call victims in an organization about an issue and request a call back. The number given is either staffed by accomplices or answered by a legitimate-sounding automated system. Victims are prompted to release potentially damaging information, about themselves in the case of identity theft, or about their company.

In another example, an attacker can leverage automated phone answering systems which allow the caller to input the first few letters of the last name of a contact to reach their extension. The attacker can try multiple extensions and stumble on some that reveal details regarding a person's position, title, or office status (e.g., "I'm out on vacation until April 4"). A social engineer can leverage that information and call other employees to gain additional information or access.

It should also be noted that with the proliferation of Voice over Internet Protocol (IP) (VoIP) there is now Voice Phishing, which is propagating a vishing-phishing attack. Since the call is routed over IP, the call number can be changed easily. This is analogous to how phishing Websites operate. Since this is primarily a voice-based social-engineering attack, it is considered a voice attack instead of a Web attack.

VoIP has also internationalized the occurrence of such frauds. In an incident in 2013, criminals thought to be in Africa who were advertising nonexistent free kittens in thousands of pages on the Web used VoIP to contact victims in an attempt to fool well-meaning cat-lovers into paying for imaginary air-shipments of imaginary cats.[36]

19.3.5.6 Social Media. Social media have exploded in popularity in recent years, and so have the opportunities for the social engineer to victimize that frequent these outlets.[37]

Useful details may be easily harvested just from searching the victim's social media page, or the pages of relatives who include the victim on their pages. Data harvested could include things such as:

- Event dates
- Friends/relatives/colleagues of interest

- Pictures
- Contact information
- Locations frequented

In addition to strictly harvesting data, a social engineer could use impersonation to befriend the individual, so that details hidden from public view are viewable. This may allow the attacker to target his attack to his needs more effectively.

Of particular concern is the LinkedIn network, which reached 187 million members worldwide by the end of 2012; 57 percent were located in the United States and 63 percent elsewhere.[38] LinkedIn makes it possible to track job titles, job history, technical background, education, and linkages to other people in any given organization—ideal information for social engineers to commandeer. Although details are supposed to be kept away from people who have not linked to each other, some LinkedIn users routinely accept link requests from anyone who sends them an invitation—even though LinkedIn's User Agreement explicitly states that members must not "Invite people you do not know to join your network."[39] Criminals have demonstrated their interest in LinkedIn contacts by selling catalogs of information gleaned on this network.[40]

19.3.6 Reverse Social Engineering. Reverse social engineering is an effective attack usually executed by an experienced social engineer. A reverse social-engineering attack has three distinct parts. First, a social engineer will create a problem, for example, a user ID issue. Second, the social engineer will publicize that they are the only person capable of fixing the issue. In the final part of the attack, the social engineer will assist the victim and "fix" the issue. It is during the third segment of the attack that a social engineer will gather information. The success rate of reverse social-engineering attacks tends to be high simply because the victim is satisfied that the fabricated problem is fixed.[41]

For example, the attacker may change the name of a file (the problem). The victim searches for and cannot find the file. The attacker will announce that they have been able to retrieve lost information but will require a user ID and password to gain access to the system (the publicity). The victim, flustered by the thought of losing an important document, will divulge the information. Finally, the attacker will "find" the missing file (the fix). The victim, being pleased that the file is returned, will forget having let out access credentials to the system.[42]

19.4 THE PSYCHOLOGY OF SOCIAL ENGINEERING

19.4.1 Introduction. Well-known security expert Bruce Schneier identified that behavioral economics, psychology of decision making, psychology of risk, and neuroscience can help explain why our feeling of security deviates from reality.[43] This section focuses on one of these aspects—psychology—and studies the science underlying the success of social engineering. Section 19.4.2 uses some well-established principles of psychology and social psychology to analyze social engineering from two angles, the psychological perspective of the victim and social-psychological perspective of the social engineer and victim. Section 19.4.3 explains that there is no single social-engineer stereotype. The terms used in this section can be found in undergraduate psychology and social psychology textbooks and therefore academic references have been minimized.

Social engineering succeeds because of human nature: In this section, examples of social-engineering attacks illustrate scientific terms used to characterize social engineering.

19.4.2 Psychology. A *cognitive bias* is defined as a mental error caused by humans' simplified information-processing strategies.[44] People become victims of social-engineering attacks due to cognitive tendencies, also called "heuristics" or rules of thumb, which are inherent in all humans. These cognitive tendencies are useful, mostly. Psychologist Robert Cialdini writes:

> We can't be expected to recognize and analyze all the aspects in each person, event, and situation we encounter in even one day. We haven't the time, energy, or capacity for it. Instead, we must very often use our stereotypes, our rules of thumb, to classify things according to a few key features and then to respond mindlessly when one or another of these trigger features is present.[45]

Although cognitive biases are found in everyone, the universal presence of cognitive biases does not imply that they are impossible to counter.

Following are some cognitive biases that can explain why people fall prey to social-engineering attacks:

- **Choice-supportive bias**. People tend to remember an option they chose that had more positive aspects than negative aspects.[46] information technology (IT) helpdesk operators may provide employee names, extensions, or both without verifying the identity of callers. Helpdesk operators remember this practice as being good because most callers are genuine and the callers thank them for providing the information. Social engineers can masquerade as genuine callers to exploit helpdesk operators' choice-support bias. According to an article on security analysis site SecurityFocus, a demonstration of this attack by the Computer Security Institute actually succeeded.[47]
- **Confirmation bias**. People tend to collect and interpret evidence in a way that confirms their conceptions.[48] If an organization has a contract with a custodial service and employees see custodians all wearing the same uniform, then a social engineer wearing the uniform may not be challenged to identify himself because of the employees' confirmation bias. The movie *Ocean's Eleven* contains many scenes in which the owner and employees of three Las Vegas casinos are conned by picaresque protagonists dressed to blend into their roles and environment.[49]
- **Exposure effect**. People tend to like things that are familiar.[50] A social engineer may call victims and explain they are performing a survey for a popular local restaurant and ask about the organization by which the victim is employed. The victims are comfortable providing that information because they are familiar with the restaurant. According to malware researcher Elodie Grandjean, some of the themes of baited emails, URLs, or instant messages used in social-engineering attacks include[51]:
  1. Pornographic links and images
  2. Using a female name in the sender field
  3. Political agendas, including solicitations for contributions in the name of a popular candidate

4. Fake emails for banks, online payment services, and other financial services. These request a confirmation or an update of login credentials or credit card information.
5. Threatening emails, mentioning jail sanctions or jury-duty procedures
6. Free games and screensavers containing a Trojan, or free antispyware tools, which are often rogue programs themselves
7. Big events, such as sports, extreme weather disaster, or urgent news
8. Celebrity names and reports on their adventures and misbehavior
9. Potentially trusted or secret relationships such as affiliation with a 12-point program
10. Social networking Websites, fake friends, school classmates or relatives, and secret lovers

- **Anchoring**. People tend to focus on one trait when making decisions.[52] If a social engineer has a soothing voice, the victim may focus on that attribute versus the questions being asked. Security expert Winn Schwartau describes a successful social-engineering attack where some employees of a New York financial services company fell for a bogus letter written on company letterhead claiming to be from the information security department; 28 percent of the 1,200 employees wrote back with their personal details.[53] The employees had focused, to their detriment, too much on the letterhead and had ignored other aspects of the request, such as its suspiciousness, graveness, and suddenness.

19.4.3 Social Psychology. Social psychologists define the *schema* as the inherent picture of reality used by humans to make judgments and decisions. From the perspective of social psychology, social engineers exploit the fact that most peoples' schema includes rules to be trustful of other people and their intentions. People are taught from the very beginning of their socialization that being nice to others is a good thing. In the context of information security, peoples' tendency to blindly trust others can spell disaster.

Below is a list of common errors that people make, and examples of how social engineers will exploit those mistakes to attack an organization.

- **Fundamental attribution error**: In this common error, people assume that the behaviors of others reflect stable, internal characteristics. Someone committing the fundamental attribution error might see a colleague in a bad mood and think, "She is *always* moody." In reality, the colleague might be pleasant in general but be suffering a headache at the time. Social engineers will act pleasant and charming to lead victims to commit the fundamental attribution error, to be impressed that the attackers are nice people *in general* and so to help them.
- **Salience effect**: Given a group of individuals, people tend to guess that the most or least influential person is the one who stands out the most. For example, from a group of ten people, nine of whom are six feet tall and one who is five feet tall, if asked to guess who the most intelligent person in the group is, an observer might say that it is the five-foot-tall person. Social engineers attempt to blend into their victim's environment to take advantage of the salience effect. They are acutely aware of company lingo, events, and regional accents.

- **Conformity, compliance, and obedience**: People respond to the social pressures of conformity, compliance, and obedience by adjusting their behaviors. A social engineer impersonating a high-powered executive demanding admittance into the company premises may persuade a new security guard with the weight of his assumed authority. The authority figure's promise of reward or threat of punishment may further influence the security guard's decision to carry out the request of the attacker.

For more details of the implications of social psychology for information assurance, see Chapter 50 in this *Handbook*.

19.4.4 The Social Engineer Profile. The profile of a social engineer is not that of the stereotypical computer hacker often portrayed in the movies or on television.[54] Social engineers are most likely not going to be loner teenagers who spend all their time with their computers in a dark basement. A social engineer is often outgoing, confident, an excellent communicator, and well-educated. Social engineers may use their own personality, or adopt a persona that greatly differs from their normal personality—oftentimes a personality they have developed over months or even years. To achieve their goals, they will blend into the environment. Social engineers seek to be unnoticeable, unremarkable. They will dress according to the dress code of the environment in which they operate. Interestingly, social engineers may be excellent actors, able to think on their feet and to adapt quickly to changing conditions. The attacker's confidence will often mask any nervousness or tension during the social-engineering attempt, which may lead to unwarranted credibility.

Social engineers may also exhibit a dark side. Attackers may have very little regard for the consequences of their actions on the victim. Even though the attackers may appear very polite or congenial towards their victims, they actually care very little about the victim or the people utilized to achieve a goal. The victim and others are simply a means to an end; they are only part of the social-engineering attack tool. The social engineer's motivations may vary widely and range from personal financial gain to revenge. There may also be significant external pressure on the attacker from acquaintances or organized crime syndicates.[55]

For more details of the psychology of computer criminals, see Chapters 12 and 13 in this *Handbook*.

19.5 DANGERS OF SOCIAL ENGINEERING AND ITS IMPACT ON BUSINESSES. The ultimate goals of the attacker may be limited to theft but may also include disruption of business or even destruction of a business. An organization must evaluate the potential impact of even a seemingly minor social-engineering attempt.

19.5.1 Consequences. The consequences of a successful social-engineering attack are almost innumerable. Anything one can dream up as an attack vector has probably been attempted by an attacker somewhere. Much like disaster recovery planning, when trying to quantify and understand the impact of social-engineering attacks, all possibilities must be put on the table. Something as seemingly minor as an internal memo that has not been properly destroyed could have the potential to bankrupt an organization. As a result, a simple social-engineering attack can lead to a more complex attack which morphs into a major information security breakdown simply based on the information contained in one misplaced memo.

The danger is especially high for publicly traded companies that can lose value because of a loss of confidence from investors.[56] Many companies, within the last few years, have come under financial duress because of a security breach or lapse that drew considerable attention from the press. Social engineering was likely part of these security incidents. Many organizations are required by law to have safeguards in place when it comes to data security; more and more of these requirements require security controls that will aid the organization in defending against social-engineering and low-tech attacks.

Another serious consequence of a successful social-engineering attack is the uncertainty that follows and the difficulty of investigating such an attack. Organizations may never fully discover to what extent a social engineer was able to infiltrate the organization. There are so many different vectors and possibilities of intrusion that it may be impossible to fully understand exactly what or who was compromised during the attack. An especially dangerous and frustrating situation is one where an attacker has an insider or accomplice within the organization. Some organizations have never identified these individuals within the organization. This uncertainty is difficult to recover from and defend against in the future. It is also another situation where a company may have substantial credibility problems and a loss of confidence from its shareholders after an attack if such details are disclosed. This exact situation has also led to a tug-of-war between regulators and organizations about disclosure laws. These laws vary widely state to state and between countries. Businesses fight to protect themselves from having to do extensive disclosure for the very reason of protecting the company's image and financial value.

19.5.2 Case Studies—Examples from Business. There are many well-known examples from real-life scenarios from organizations that can demonstrate the effectiveness of social engineering. Here are some well-known examples without any specifics that would embarrass the involved parties. Social-engineering attacks are real—they are not simply computer-security theory.[57]

Case 1: One very well-known social-engineering attack in the business world is piggybacking. In this type of social engineering, an attacker will depend on an individual's sense of courtesy. Most people remember from their early school days that they are taught to hold a door open for someone who is behind him or her. This courtesy is often extended in the workplace, including secure areas such as a datacenter. A potential attacker may attempt to enter a datacenter without proper credentials by following closely behind someone else who is entering a datacenter with proper identification and authorization. The authorized individual will probably exhibit courtesy and hold the door open for the piggybacker, even if that person's identity is unknown. Holding something innocuous often creates confidence.

Result: In this example, the result is that an attacker who is not authorized has gained access to a secure facility by relying on another individual's sense of courtesy. It can be very difficult to demand that employees refrain from allowing piggybacking, but a policy must be in place to demand exactly that. All persons entering a secure facility should be required to fully use identification and authorization mechanisms every time.

Case 2: Another example that has been carried out in different variations involves an attacker using several social-engineering techniques to take advantage of

several coinciding events to exploit data, information, or equipment from an organization. The attacker would often make several telephone calls or send several emails to find a specific date where a company official, perhaps the CFO or director of technology, is out of the office. The attacker would then show up at the organization making claims that the company official authorized the attacker to take a certain computer from the company's premises. Usually the attacker will try to show a sense of urgency and extend a very confident display of authority. Many times the employee will cave into the attacker's demands without checking the story, and now a very important computer has been taken from the company. (A very commonly used twist on this case, and one that is becoming more and more common, is to install malware, remote backdoor software, or a keylogger on the computer instead of physically removing it.)

Result: In this case, the organization has lost control and ownership of a computer and/or its data. If the computer does not have any solid safeguards, such as data encryption mechanisms, the data and the computer could be used for any number of destructive activities. If information about the incident is made public, great damage to the organization's reputation can be done. The data harvested from the computer could also be sold to competitors or used as part of a blackmail scheme.

19.5.3 Success Rate. Although there are few accurate statistics of the success rate of social engineering, as is the case in many areas of information security, most experts believe the rate to be high. If history has anything to teach the security community through example, social engineering will continue to be a powerful and successful tool for criminals. Few organizations are immune to social engineering, no matter their industry or product. If even well-trained military personnel are vulnerable, every organization must take the success rate seriously.

The high success rate must also stress the importance of continued education of employees and reevaluation of the organization's efforts to combat social engineering and protect all data assets. This is an area where many organizations, of all sizes, continue to allocate too few resources. The frequency of social-engineering attempts also dictates the need for proper, efficient, and swift reporting of suspicious activity targeting individuals or the organization. It is very difficult to defend against social-engineering attempts if the organization does not know it is under attack or know where or how to report an incident. The probability of a successful attack can be substantially reduced with properly trained, supported, and motivated employees.

19.5.4 Small Business verses Large Organizations. The impact and dangers of social-engineering and low-tech attacks vary widely between small businesses and large corporations. As discussed above, the consequences can be potentially serious, all the way up to the collapse of the organization. Small businesses are often much less prepared operationally and financially and much less equipped to survive a serious breach of security. Conversely, and perhaps shocking to many, small businesses may have the upper hand over large organizations because of a substantially smaller workforce and collapsed management structure. It is much easier to communicate and engage everyone within a small company when an attack is attempted or carried out. Small businesses also have an advantage of a much smaller workforce to train; this results in much better prepared employees. Small business employees are probably

more likely to identify people who do not belong in the organization or should not be asking for access to sensitive areas or data. They may also be more likely to deny access or question someone whose story does not seem likely or is suspicious.

Large organizations can be mired in bureaucracy, ineffective management, or overly complicated reporting procedures. An attacker could carry out an entire plan before the security team in a large company would be alerted to a social-engineering attempt. Many times an individual is less willing to challenge the credentials of a stranger in a large organization. This is especially true where there are employees that feel they may be punished for questioning someone of a higher rank or preventing someone else from doing his or her job. The employee might be more likely to just let the attacker pass unquestioned rather than risk possible negative ramifications or scrutiny.

Bottom line, all organizations must be on the lookout for this type of attack. Criminals do not always choose easy or obvious targets. Any business, small or large, family-owned or corporate conglomerate, may be a target of an attack that utilizes social engineering.

19.5.5 Trends. While it is important for any information security manager to always keep a skeptical eye toward statistics, it is equally important to keep abreast of security threat trends. Social engineering is no exception. It may be difficult for any survey or poll to gather facts about how many attempts were made in any given year or how many were successful. Many social-engineering attempts are probably never detected, and even less are reported or admitted to on surveys. When it comes to such a powerful and successful attack mechanism, assume the worst: It is increasingly being used by criminals even when organizations know about this type of attack and even train employees to defend against it. Criminals create new forms and tactics every day, and people will probably continue to fall for these tactics.

19.6 DETECTION. Detection of social-engineering and low-tech attacks can be difficult. The nature of most types of social-engineering attacks is to take advantage of people's willingness to trust and help, which allows the perpetrator to circumvent technical controls. In many cases, the detection relies on people's ability to recognize a potential attack, including suspicious activity and respond appropriately. Further complicating detection is the potential that a social-engineering attack may not be a single occurrence, but many smaller events culminating in the attackers accessing restricted resources, hide nefarious activities, or stealing proprietary information. Using slow or multistep social-engineering attacks is even more difficult to detect and potentially more damaging. More recently, successful social-engineering attacks through the targeted email campaigns to senior level personnel have been widely reported. One penetration-testing expert estimates that social-engineering tactics account for less than 20 percent of the time spent in an attack; the rest of the time is spent in technical exploitation of the information originally gathered through deceit.[58]

There are three main avenues of detection for social-engineering attacks: people, audit controls, and technology.

19.6.1 People. Since people are the vector for the attack, they are, in general, the first line of defense. Organizations need to provide employees with the resources and continuous education to help discern a potential attack from a legitimate request. All areas of an organization need the training. Social engineers are using all available resources; there is no department or person that is immune from a potential attack. In addition, organizations need to provide information on what to do during and

immediately after an attack. For example, during a phone attack, employees should be trained to remember as many details as possible. Items that the employee should try to remember include:

- Was the attacker male or female?
- Was a caller ID displayed?
- Was there noise in the background?
- Did he or she have an accent?
- What questions were asked?
- What answers were provided?

Employees should also be aware of people asking many questions, some of which may not make sense. In addition, callers or emails requesting names of managers or IT personnel should prompt a phone call to the IT security or investigation department. Organizational culture is also another pivotal piece to help defend against social-engineering attacks. The culture must stress and *reward* proper verification before providing information. Upper management can support a strong culture by praising employees who ask *them* to demonstrate their authenticity, rather than expressing irritation that anyone would have the nerve to do so.

All businesses, including security firms, must adhere to that mantra.[59] Unfortunately, too often VIPs or senior leaders do not have to follow the same verification process to change passwords or receive information. Finally, employees should be aware of their own policies and business practices. For example, all employees should understand why it is inappropriate that a caller would request a password change via phone or text message.[60]

Organizations also need to provide clear information to their associates as to what to do when confronted with possible social engineering. For example, who should be notified during the actual event, immediately afterward, or both (depending on when the employee recognizes the attack). To help in the notification process, an organization can create an incident-notification information system and disseminate the information across the organization so an employee can immediately and easily determine whom to contact for help.

19.6.2 Audit Controls and Event Logging. Auditing and logging email, Internet content, systems logins, and systems changes of an organization can be used to detect social-engineering attacks. If the review of events does not happen in real time, there will be a delay from the time of the attack and identification of incident. In the instances where notification is not real time, forensic examiners can use the audit information to help piece together the attacks and determine root cause. Awareness teams can also use the log results to help devise additional training for the targeted groups. During real-time auditing, organizations can immediately enact their incident management plans to help limit the potential damage. Modern anti-malware and anti-phishing tools can provide not only defenses against such attacks but can also keep detailed records for analysis.

19.6.3 Technology for Detection. Organizations can implement technology to help limit social-engineering attacks. Content-filtering software can limit email and Website traffic; any identified malware should be blocked or stripped from email.

Email monitoring tools can be used in bidirectional mode to inspect content in both directions. In addition, email-monitoring and content-filtering mechanisms can be used to scan for keywords or phrases that may trigger early warning signals. Any suspicious traffic, attachment, or email should be quarantined and reviewed prior to delivery. By scanning and blocking content, organizations may reduce the number of suspicious emails entering their networks and reduce the number of suspicious Websites that employees visit. Email services and anti-malware products have been improving their spam-blocking capabilities so that even phishing scams are being trapped successfully.

Advances in technology research will provide additional protection from social-engineering attacks, including the development of the Social Engineering Defense Architecture (SEDA). This architecture attempts to detect social-engineering attacks over the phone by identifying a legitimate employee versus a social engineer. The system is referred to as a text-independent voice signature authentication system. The system uses voice recognition technology, which would reduce the risk of a successful helpdesk attack. In addition to detecting an unauthorized caller, it can detect an insider masquerading as an employee with a higher security classification. The logging included in the architecture would aid a forensic examiner during an investigation.[61]

Risk-based authentication methods can also help limit potential social-engineering attacks; for example, if a customer or user always calls from a certain area code or phone number and a potential attacker calls from a different area code, systems can identify this and prompt for additional information prior to authenticating the individual.[62]

Attackers are becoming more sophisticated. Recent phishing attacks are targeting systems, network, and security professionals. In these types of attacks, a phishing email is sent to an organization from a "customer" informing it of a phishing site attempting to steal customer information. Security personnel respond to the email and investigate the site. The site installs malware that allows the attacker to remotely control the machine. This change in attack methodology requires security personnel to be even more suspicious of any seemingly innocuous email telling the organization that there are phishing sites targeting that company.

19.7 RESPONSE. Responding to social-engineering and other low-tech attacks should fit into an organization's incident-management process and response. As with all incident management plans, the responses should be well-defined, communicated, and tested. It behooves an organization to plan for the inevitable, especially as network and physical attacks are increasingly using multiple vectors in order to be successful.

See Chapter 53 in this *Handbook* for additional information on monitoring and controlling systems and see Chapter 56 for information on computer-security incident-response teams.

19.8 DEFENSE AND MITIGATION. The prevention of social-engineering attacks should be multifaceted, repeatable, and part of an organization's defense-in-depth strategy. Since the very nature of the attack is to bypass or circumvent technical defenses by using people's good nature and willingness to trust other human beings, the steps in preventing such attacks should focus on distinct areas, such as policy, training and awareness, technology, and physical defenses.

Focusing on all areas will help mitigate the threat of a successful social-engineering attack. Each area should have a regular process for review and auditing. Well-written policies provide the baseline of behavior that is acceptable and the potential consequences if they are not followed, but they must be accessible to employees. In the case of training, it should be integrated into the organization's overall security awareness

program, included in all employee's evaluations, and tied to bonus pay and compensation. Organizations must train their employees in acceptable behavior and provide policies and the tools to identify and report potential social-engineering attacks.

Physical defenses can potentially block an intruder from entering a locked building, but people's sense of camaraderie does make the practices such as piggybacking or tailgating relatively easy. A recent study indicated that smokers returning to a building sometimes allow nonemployees into a secure location. Technological advances are also important but cannot be relied on as the only method of defense.

19.8.1 Training and Awareness. Organizations need to provide tools and knowledge to employees to help identify potential attacks and react to suspected attacks. Providing awareness is a continual process and should not be only for new hires. Training and awareness, when possible, should include real-life examples so employees can relate to the issue and understand the level of trust that is implied through their access to the facility and the data used for their positions. Employees should understand the responsibilities the company bestows on them. Essentially, employees need to be retrained that it is acceptable to ask why certain information is being requested or to see the badge of a person behind them.

A basic awareness program could include posters, email communications, and laminated instruction cards (hard cards) containing emergency contact numbers or other information. More mature awareness programs can include videos or brown-bag informational lunches. Whenever possible, teaching employees to defend against social-engineering attacks should be a live presentation with real-world examples. Ideally, the presentation should be tailored specially for each audience. For example, if the audience is helpdesk personnel, examples of potential social-engineering attacks should be described or demonstrated and discussed.

Employee training can also include instruction to keep file cabinets locked when not in use, lock workstations, and use cable locks, and contain instructions on how to create and remember good passwords.

The backbone of awareness training is a well-defined and executed information-security policy program. Please see Chapter 44 in this *Handbook* for additional information regarding information-security policy development, and Chapter 49 for additional information regarding awareness programs.

19.8.2 Technology for Prevention. Technology is emerging as a defense against certain types of social-engineering attacks. The technology enables organizations to identify some social-engineering attacks without relying on employees. This proactive identification enables an organization to mitigate the risk. Technology should comprise only one layer of defense and not be relied on as the sole defense.

Technologies such as content-monitoring systems for both email and Web content can help identify phishing attacks. In addition, organizations can install timely security patches and use up-to-date antivirus and antispyware software to help mitigate the risk of viruses, Trojans, and worms. Most versions of browsers and browser plug-ins are allowing users to evaluate the trustworthiness of Websites.[63]

Ideally, employees should be prevented from downloading and installing unapproved software. However, organizations can employ inventory systems or other methodologies to detect illegal programs on a network. Certain types of systems can prevent malware-infected machines from entering the network.

Desktop and laptop configuration changes can be made to reduce the risk of a successful attack; changes include disabling pop-up windows in browsers, disallowing

the automatic installation of Active-X controls, multiple versions of Java, limiting the types of cookies that Websites can place on local machines, using automatic password-protected screen savers, and finally, using email certificates for authentication.

In the same regard, organizations should review technology processes and verify that they are not inadvertently supplying information to potential social engineers. For example, metadata in documents should be removed before being accessible to outsiders, and regular Web searches should be conducted to ensure former or current employees are not posting information on the Internet.

19.8.3 Physical Security. Physical security mechanisms can reduce the risk of a successful social-engineering attack. All employees should have identification cards that they are required to display at all times. Secured areas within an organization should be locked, have limited access, and be monitored for noncompliance. Door alarms that can detect tailgating or piggybacking can be installed in areas. Cameras or other closed circuit monitoring technology can thwart potential intruders. Security personnel should watch all facility access points. All office doors, desks, file cabinets, and other storage devices should have keys and remain locked when not being accessed. Dumpsters® or recycling bins should also have locks that would prevent the removal of documents meant for shredding or incineration.

- Desktops, laptops, and other computer hardware should be physically locked in place. Users should be required to have strong passwords, and in the case of laptops or desktops, have automatic screen savers that require a password to unlock. All magnetic media need to have secure storage.
- Most organizations have a certain percentage of mobile workforces. Special training should be provided to them to prevent the loss of equipment or information. Training should include information such as:
  - Laptops should remain with the traveler and not checked in luggage.
  - Laptops should be locked in a safe or to a secure surface at all times.
  - Peripheral devices such as USB drives and handheld devices should have strong passwords.
  - Conversations involving confidential information should be prohibited in public.
  - Travelers should be aware of their surroundings, and special considerations regarding electronic communication and equipment usage should be considered while not in the United States and Europe.

Please see Chapters 22 and 23 of this *Handbook* for more information regarding physical security.

19.9 CONCLUDING REMARKS. Social-engineering attacks are unlike technological computer attacks. The vector of the attack is human, the nature of the attack is to circumvent controls, and the success of the attack depends on people's willingness to trust others. Simply being polite and holding a door open for a stranger can be bad for an organization, if the stranger is a social engineer.

Social-engineering attacks are not new in society; they have been used for millennia. Social engineers use many different methods to execute social-engineering and low-tech attacks. These methods could involve human contact, no human contact, or a combination of technology and social-engineering tactics.

Psychologists and social psychologists have offered a number of reasons for the success of social engineering. They theorize that human nature allows attackers to trick or con reasonable people into divulging information. A social engineer's profile does not fit into a single model, and their attacks are difficult to detect. Social-engineering attacks have grown more prolific, effective, and dangerous to both organizations and individuals.

Even though social-engineering attacks are difficult to defend against, there are technical, process, and personal defenses that an organization can adopt to prevent them and minimize their effects. As with all information security issues, a defense-in-depth strategy can help mitigate the risks associated with social engineering.

Between July and August 2011, Dimensional Research surveyed 835 IT professionals in the Western economies. Some of their findings reinforce the details in this chapter[64]:

- The threat of social engineering is real.
- Financial gains are the primary motivation of social engineering.
- Social-engineering attacks are costly, especially in large organizations.
- [There is a] Lack of proactive training to prevent social-engineering attacks.

Although the focus of information assurance practice seems to be on technical attacks, social-engineering and low-tech attacks will continue to remain relevant and dangerous threats, worthy of attention and mitigation.

19.10 NOTES

1. M. E. Kabay, "Social Engineering Simulations," Network World Security Strategies, December 18, 2000, www.mekabay.com/nwss/059_social_engineering_v03.pdf
2. Ira Winkler, *Spies Among Us* (Indianapolis, IN: Wiley Publishing, Inc., 2005): 111.
3. L. Giles, trans., *Sun Tzu on the Art of War: The Oldest Military Treatise in the World*, www.chinapage.com/sunzi-e.html
4. Charles R. King, *Learning Secrets: Intelligence in the Second Punic War*, Kindle ed., Amazon Digital Services, Inc., 2012, www.amazon.com/Learning-Secrets-Intelligence-Second-ebook/dp/B0082PZIYI
5. Ben Macintyre, *Double Cross: The True Story of the D-Day Spies* (New York: Crown, 2012).
6. Edward Lucas, *Deception: The Untold Story of East-West Espionage Today* (New York: Walker & Company, 2012).
7. Steven Spielberg (dir.), *Catch Me If You Can*, 2002, www.imdb.com/title/tt0264464
8. Frank W. Abagnale, Website,www.abagnale.com/index2.asp
9. Stacy Cowley, "How a lying 'social engineer' hacked Wal-Mart," *CNN-Money*, August 8, 2012, http://money.cnn.com/2012/08/07/technology/walmart-hack-defcon/index.htm
10. For more information on this documentary, see www.hbo.com/docs/programs/hackingdemocracy/index.html
11. Andy O'Donnell, "Caller ID Spoofing," *About.com: Internet/Network Security | Cybercrime*, May 2, 2012, http://netsecurity.about.com/od/securityadvisorie1/a/Caller-Id-Spoofing.htm

12. Joan Goodchild, "4 Reasons Why Executives are the Easiest Social Engineering Targets," CSOonline.com, July 14, 2012, www.csoonline.com/article/599456/4-reasons-why-executives-are-the-easiest-social-engineering-targets
13. "Social Engineering," AuditMyPC.com, www.auditmypc.com/social-engineering.asp
14. Dan Arnell, "Payroll Giant Gives Scammer Personal Data of Hundreds of Thousands of Investors," *ABC News*, June 26, 2006, http://abcnews.go.com/Technology/story?id=2160425&page=1
15. M. E. Kabay, personal communication, 2013.
16. "Seduce," *The Free Dictionary*, www.thefreedictionary.com/seduce
17. Federal Bureau of Investigation, "Citadel Malware Continues to Deliver Reveton Ransomware in Attempts to Extort Money," FBI New E-Scams & Warnings Website, November 30, 2013, www.fbi.gov/scams-safety/e-scams
18. M. E. Kabay, "Office Supply Scams," M. E. Kabay Website, 2011, www.mekabay.com/nwss/891_office_supply_scams.pdf
19. Bureau of Consumer Protection, "Avoiding Office Supply Scams," Federal Trade Commission, March 2000, http://business.ftc.gov/documents/bus24-avoiding-office-supply-scams
20. Ajay Gupta, "The Art of Social Engineering," Addison-Wesley, August 23, 2002, www.informit.com/articles/article.aspx?p=28802
21. Michael Winter, "Confetti from police files tossed at Macy's parade," *USA TODAY*, November 26, 2012, www.usatoday.com/story/news/nation/2012/11/26/thanksgiving-macys-parade-confidential-police-confetti/1728375/
22. M. Horowitz, "Defending against Malicious CDs and USB Flash Drives," *Computerworld | Defensive Computing* Blog, July 3, 2011, http://blogs.computerworld.com/18560/defending_against_malicious_cds_and_usb_flash_drives
23. Microsoft Support. "How To Minimize Metadata in Office Documents," Microsoft Corporation, January 24, 2007, http://support.microsoft.com/default.aspx?scid=kb;EN-US;Q223396
24. The 419 Coalition, "The Nigerian Scam (419 Advance Fee Fraud) Defined," Nigeria—The 419 Coalition Website, accessed March 30, 2007, http://home.rica.net/alphae/419coal
25. *Webopedia*, s.v. "Phishing," accessed March 30, 2007, www.webopedia.com/TERM/P/phishing.html
26. Robert Richardson, *2010/2011 Computer Crime and Security Survey*, Information Week, Computer Security Institute, June 6, 2011, http://reports.informationweek.com/abstract/21/7377/security/research-2010-2011-csi-survey.html
27. RSA FraudAction Research Lab, "Anatomy of an Attack," RSA Speaking of Security Blog, April 1, 2011, http://blogs.rsa.com/rivner/anatomy-of-an-attack
28. *SearchExchange*, s.v. "Spim (Instant Messaging Spam)," accessed March 30, 2007, http://searchexchange.techtarget.com/sDefinition/0,290660,sid43_gci952820,00.html
29. Will Sturgeon, "U.S. Makes First Arrest for Spim," *CNet.com*, February 21, 2005, http://news.cnet.com/U.S.-makes-first-arrest-for-spim/2100-7355_3-5584574.html
30. Linda Rosencrance, "Porn Site WhiteHouse.com Domain Name Up for Sale," *Computerworld*, February 10, 2004, www.computerworld.com/s/article/90035/Porn_site_WhiteHouse.com_domain_name_up_for_sale

31. Ryan Naraine, "Hushmail DNS Attack Blamed on Network Solutions," eWeek.com, April 29, 2005, http://www.eweek.com/c/a/Security/Hushmail-DNS-Attack-Blamed-on-Network-Solutions/
32. Steve Stasiukonis, "Social Engineering, the USB Way," Dark Reading, June 7, 2006, http://www.darkreading.com/perimeter/social-engineering-the-usb-way/208803634
33. Richardson, *2010/2011 Computer Crime and Security Survey*
34. Jonathan Ness, "More Information about the DLL Preloading Remote Attack Vector," Microsoft Security Research & Defense, August 23, 2010, http://blogs.technet.com/b/srd/archive/2010/08/23/more-information-about-dll-preloading-remote-attack-vector.aspx
35. Brian Koerner, "Vishing," About.com: Vishing, accessed March 22, 2007, en.wikipedia.org/wiki/voice_phishing
36. M. E. Kabay, "Kitty Porn," *InfoSec Perception* Website, September 28, 2012, http://resources.infosecskills.com/perception/kitty-porn
37. Patterson Clark and Robert O'Harrow, Jr., "Social Engineering: Using Social Media to Launch a Cyberattack." *The Washington Post*, September 26, 2012, www.washingtonpost.com/investigations/social-engineering-using-social-media-to-launch-a-cyberattack/2012/09/26/a282c6be-0837-11e2-a10 c-fa5a255a9258_graphic.html
38. Julie Inouye, "Oh, What a Year It's Been!" LinkedIn Blog, December 21, 2012, http://blog.linkedin.com/2012/12/21/oh-what-a-year-its-been/
39. LinkedIn, "User Agreement," last revised May 13, 2013, www.linkedin.com/static?key=user_agreement
40. Stacy Cowley, "LinkedIn Is a Hacker's Dream Tool," *CNNMoney*, March 12, 2012, http://money.cnn.com/2012/03/12/technology/linkedin-hackers/index.htm
41. Sarah Granger, "Social Engineering Fundamentals, Part I: Hacker Tactics," Symantic Website, December 18, 2001, updated November 3, 2010, www.symantec.com/connect/articles/social-engineering-fundamentals-part-i-hacker-tactics
42. Microsoft Corporation, "How to Protect Insiders from Social Engineering Threats," *Microsoft Technet*, August 18, 2006, http://technet.microsoft.com/library/Cc875841#EIXAE
43. Bruce Schneier, "The Psychology of Security," Essays and Op Eds, 2007, www.schneier.com/essay-155.html
44. Richards J. Heuer, Jr., "What Are Cognitive Biases?" in *Psychology of Intelligence Analysis*, Part 3—Cognitive Biases, 1999, https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/psychology-of-intelligence-analysis/art12.html
45. R. Cialdini, *Influence: The Psychology of Persuasion* (New York: HarperCollins, 1998).
46. Mara Mathers, Eldar Shafir, and Marcia K. Johnson, "Misremembrance of Options Past: Source Monitoring and Choice," *Psychological Science* 11, No. 2, March 2000, http://psych.princeton.edu/psychology/research/shafir/pubs/Misremembrance%202000.pdf
47. Granger, "Hacker Tactics."
48. J. St. B. T. Evans, J. L. Barston, and P. Pollard, "On the Conflict between Logic and Belief in Syllogistic Reasoning," *Memory and Cognition*,11 (1983): 295–306.

49. Steven Soderbergh (dir.), *Ocean's Eleven*, 2001, www.imdb.com/title/tt0240772
50. R. B. Zajonc, "Attitudinal Effects of Mere Exposure," *Journal of Personality and Social Psychology* 9, No. 2 (1968): 1–27.
51. E. Grandjean, "A Prime Target for Social Engineering Malware," *McAfee Security Journal*, Fall (2008): 9–12.
52. A. Tversky and D. Kahneman, *Judgment under Uncertainty: Heuristics and Biases*, *Science*, 185 (1974): 1124–1130.
53. Joan Goodchild, "Social Engineering Stories," CSO, May 24, 2010, www.csoonline.com/article/594924/social-engineering-stories
54. Tim Wilson, "Eight Faces of a Hacker," *Dark Reading*, March 29, 2007, www.darkreading.com/security/news/208804443/eight-faces-of-a-hacker.html
55. Malcolm Allen, "Social Engineering, A Means to Violate a Computer System," SANS Institute Information Security Reading Room, June 2006, www.sans.org/reading_room/whitepapers/engineering/social-engineering-means-violate-computer-system_529
56. Radha Gulati, "The Threat of Social Engineering and Your Defense Against It," SANS Institute Information Security Reading Room, 2003, www.sans.org/reading_room/whitepapers/engineering/threat-social-engineering-defense_1232
57. Granger, "Hacker Tactics."
58. Kelly Jackson Higgins, "Social Engineering Gets Smarter," *Dark Reading*, June 16, 2006, www.darkreading.com/authentication/social-engineering-gets-smarter/208803849
59. P. Roberts, "Hacker Grrl Behind HBGary Attack," *threatpost*, March, 17, 2011, http://threatpost.com/hacker-grrl-behind-hbgary-attack-031711/75037
60. Ryan Groom, "Top 5 Social Engineering Techniques," About: Business Security, accessed March 17, 2007, http://bizsecurity.about.com/od/physicalsecurity/a/topsocialengine.htm
61. Michael Hoeschele and Marcus Rogers, "Detecting Social Engineering," in *Advances in Digital Forensics: IFIP International Conference on Digital Forensics* (February 2005): 67–71.
62. RSA, "Risk-Based Authentication," Information Security Glossary, RSA Website, accessed April 12, 2012, www.rsa.com/glossary/default.asp?id=1094
63. M. J. Edwards, "IE 7.0 and Firefox 2.0 Both Have New Antiphishing Technologies," WindowsITPro, October 26, 2006, http://windowsitpro.com/networking/ie-70-and-firefox-20-both-have-new-antiphishing-technologies
64. Dimensional Research, "The Risk of Social Engineering on Information Security: A Survey of IT Professionals," Check Point Software Technologies Website, September 2011, www.checkpoint.com/press/downloads/social-engineering-survey.pdf

CHAPTER 20

SPAM, PHISHING, AND TROJANS: ATTACKS MEANT TO FOOL

Stephen Cobb

20.1 UNWANTED EMAIL AND OTHER PESTS: A SECURITY ISSUE 20·1
- 20.1.1 Common Elements 20·2
- 20.1.2 Chapter Organization 20·3

20.2 EMAIL: AN ANATOMY LESSON 20·3
- 20.2.1 Simple Mail Transport Protocol 20·3
- 20.2.2 Heads-Up 20·4

20.3 3 SPAM DEFINED 20·6
- 20.3.1 Origins and Meaning of Spam (not SPAM®) 20·7
- 20.3.2 Digging into Spam 20·8
- 20.3.3 Spam's Two-Sided Threat 20·12

20.4 FIGHTING SPAM 20·18
- 20.4.1 Enter the Spam Fighters 20·18
- 20.4.2 A Good Reputation? 20·18
- 20.4.3 Relaying Trouble 20·20
- 20.4.4 Black Holes and Block Lists 20·20
- 20.4.5 Spam Filters 20·21
- 20.4.6 Network Devices 20·24
- 20.4.7 Email Authentication 20·25
- 20.4.8 Industry Initiatives 20·26
- 20.4.9 Legal Remedies 20·26

20.5 PHISHING 20·27
- 20.5.1 What Phish Look Like 20·27
- 20.5.2 Growth and Extent of Phishing 20·29
- 20.5.3 Where Is the Threat? 20·30
- 20.5.4 Fighting Phishing 20·31

20.6 TROJAN CODE 20·31
- 20.6.1 Classic and Recent Trojans 20·32
- 20.6.2 Basic Anti-Trojan Tactics 20·34
- 20.6.3 Lockdown and Quarantine 20·35

20.7 CONCLUDING REMARKS 20·36

20.8 FURTHER READING 20·36

20.9 NOTES 20·36

20.1 UNWANTED EMAIL AND OTHER PESTS: A SECURITY ISSUE. Three oddly named threats to computer security are addressed in this chapter: *spam*, *phishing*, and *Trojan code*. Spam is unsolicited commercial email. Phishing is the use of deceptive unsolicited email to obtain—to fish electronically for—confidential information. Trojan code, a term derived from the Trojan horse, is software designed to achieve unauthorized access to systems by posing as legitimate applications. In this chapter, we

outline the threats posed by spam, phishing, and Trojans, as well as the mitigation of those threats.

These threats might have strange names, but they are no strangers to those whose actions undermine the benefits of information technology. Every year, for at least the last three years, the U.S. Internal Revenue Service (IRS) has had to warn the public about email scams that take the name of the IRS in vain, attempting to defraud taxpayers of their hard-earned money by aping the agency's look and feel in email messages. The fact that the IRS states that the agency never ever uses email, text messaging, or social-media channels to communicate with taxpayers[1] is a sad comment on our society, for there is no good reason why this should be the case. After all, the agency allows electronic filing of annual tax returns, and every day hundreds of millions of people around the world do their banking and bill paying online, in relative safety. The technology exists to make email safe and secure. In many ways, this chapter is a catalog of what has happened because we have not deployed the technology effectively.

20.1.1 Common Elements. Each of these threats is quite different from the other in some respects, but all three have some important elements in common; first, and most notably, they use deception. These threats prey on the gullibility of computer users and achieve their ends more readily when users are ill-trained and ill-informed (albeit aided and abetted, in some cases, by poor system design and poor management of services, such as broadband connectivity).

Second, all three attacks are enabled by system services that are widely used for legitimate purposes. Although the same might be said of computer viruses—they are code and computers are built to run code—the three threats that are the focus of this chapter typically operate at a higher level, the application layer of the Open Systems Interconnection (OSI) model (discussed in Section 6.5.1 in this *Handbook*). Indeed, this fact may contribute to the extent of their deployment—these threats can be carried out with relatively little technical ability, relying more on the skills associated with social engineering than with coding. For example, anyone with an Internet connection and an email program can send spam. That spam can spread a ready-made Trojan. Using a toolkit full of scripts, you can add a Website with an input form to your portfolio and go phishing for personal data to collect and abuse.

Another reason for considering these three phenomena together is the fact that they often are combined in real-world exploits. The same mass emailing techniques used to send spam may be employed to send out messages that spread Trojan code. As described in Chapter 15 in this *Handbook*, Trojan code may be used to aid phishing operations. Systems compromised by Trojan code may be used for spamming, and so on. What we see in these attacks today are the very harmful and costly result of combining relatively simple strategies and techniques with standards of behavior that range from the foolish and irresponsible to the unabashedly criminal.

These three threats also share the distinction of having been underestimated when they first emerged. All three have evolved and expanded in the 21st century. For example, for 2012, Symantec reported that Android-operating-system threats (virtually unknown a decade earlier) against mobile devices increased drastically between January 2010 and the end of 2012, from fewer than 100 unique variants in about 10 families to around 4,500 variants in about 170 families.[2] In addition, an increasing number of phishing scams in 2012 were being carried out through social media—a term not even mentioned in the corresponding report for the first half of 2006.[3] (For more information about viruses and worms, spyware, rootkits, and other malware, see Chapter 16 in this *Handbook.*)

One other factor unites these three threats: their association with the emergence of financial gain as a primary motivator for writing malicious code and abusing Internet connectivity (as discussed in Chapter 2 in this *Handbook*). The hope of making money is the primary driver of spam. Phishing is done to facilitate fraud for gain through theft of personal data and credentials, either for use by the thief or through resale in the underground economy. Trojan code is used to advance the goals of both spammers and perpetrators of phishing scams. In short, all three constitute a very real threat to computer security, the well-being of any computer-using organization, and the online economy.

20.1.2 Chapter Organization. After a brief email anatomy lesson, each of the three threats addressed by this chapter is examined in turn. The email anatomy lesson is provided because email plays such a central role in these threats, enabling spam and phishing and the spread of Trojan code. Responses to the three threats are discussed with respect to each other, along with consideration of broader responses that may be used to mitigate all three.

For an introduction to the general principles of social engineering, see Chapter 19 in this *Handbook*.

20.2 EMAIL: AN ANATOMY LESSON. Email plays a role in numerous threats to information and information systems. Not only does it enable spam and phishing, it is used to spread Trojan code, viruses, and worms. A basic understanding of how email works will help to understand these threats and the various countermeasures that have been developed.

20.2.1 Simple Mail Transport Protocol. All email transmitted across the Internet is sent using an agreed-on industry standard: the Simple Mail Transport Protocol (SMTP). Any server that speaks SMTP is able to act as a Mail Transfer Agent (MTA) and send mail to, and receive mail from, any other server that speaks SMTP. To understand how simple SMTP really is, Exhibit 20.1 presents an example of an SMTP transaction.[4] The text that follows represents the *actual* data being sent and received by email servers (as viewed in a telnet session, with the words in CAPS, such as HELO and DATA, being the SMTP commands defined in the relevant standards, and the numbers being the standard telnet responses).

Developed at a time when computing resources were relatively expensive, and designed to operate even when a server was processing a dozen or more message connections per second, the SMTP conversation was kept very simple in order to be very brief. However, that simplicity is both a blessing and a curse. As you can see from the example, only two pieces of identity information are received before the mail is delivered: the identity of the sending server, in this case example.com, and the From address, in this case foo@example.com. SMTP has no process for verifying the validity of those identity assertions, so both of those identifiers can be trivially falsified. The remaining contents of the email, including the subject and other header information, are transmitted in the data block and are not considered a meaningful part of the SMTP conversation. In other words, no SMTP mechanism exists to verify assertions such as "this message is from your bank and concerns your account" or "this message contains the tracking number for your online order" or "here is the investment newsletter that you requested."

As described in more detail later, some email services do perform whitelist or blacklist look-ups on the Internet Protocol (IP) address of the sending server during

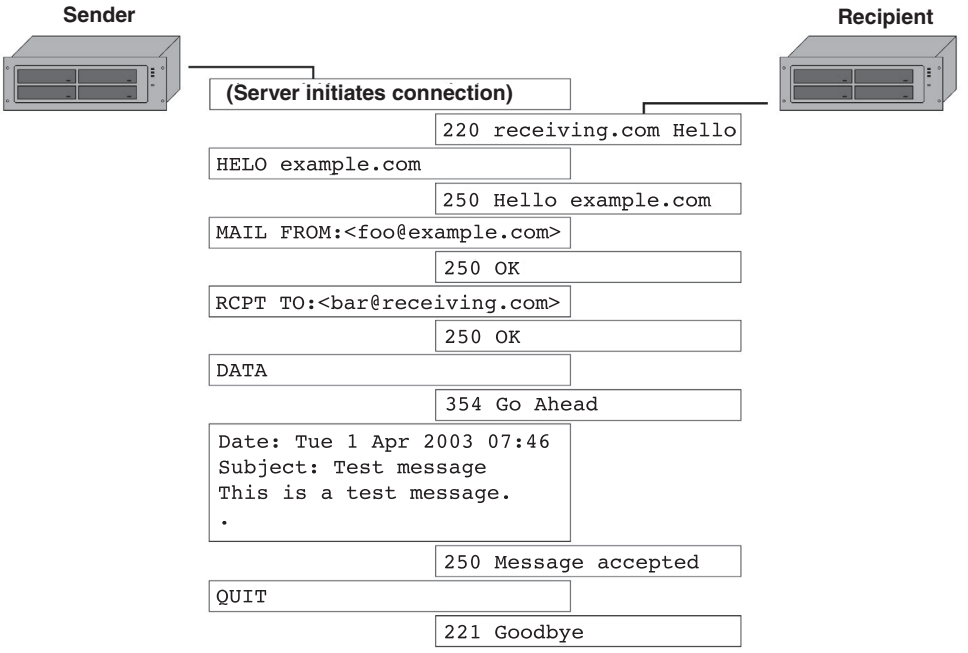


Exhibit 20.1 Basic Email Protocol

the SMTP conversation, but those inquiries can dramatically slow mail processing, requiring extra capacity to offset the loss of efficiency. A whitelist identifies email senders that are trusted; a blacklist identifies those that are not trusted. Maintenance of whitelists can be time consuming, and blacklists have a long history of inaccuracies and legal disputes. In short, the need for speed creates a system in which there are virtually no technical consequences for misrepresentations in mail delivery. And this is precisely why spammers have been, and continue to be, incredibly effective at getting unwanted email delivered.

SMTP is, as Winston Churchill might have put it, the worst way of doing email, except for all the others that have been tried. The reality is that SMTP works reliably and has been widely implemented. To supplant SMTP with anything better would mean a wholesale redesign of the entire global email infrastructure, a task that few in the industry have been willing to undertake. Some people have endeavored to develop solutions that can ride atop the existing SMTP infrastructure, allowing SMTP to continue functioning efficiently while giving those who use it the option of engaging more robust features that help differentiate legitimate mail from spam. There is more on these solutions later in the chapter.

20.2.2 Heads-Up. Email cannot be delivered without something called a header, and every email message has one, a section of the message that is not always displayed in the recipient's email program but is there nonetheless, describing where the message came from, how it was addressed, and how it was delivered. Examining the header can tell you a lot about a message. Consider how a message appears in Microsoft Outlook Express, as illustrated in Exhibit 20.2.

At the top you can see the From, the To, and the Subject. For example, you can see part of a message that appears to be from zjjimwalker8467r21@lycos.com

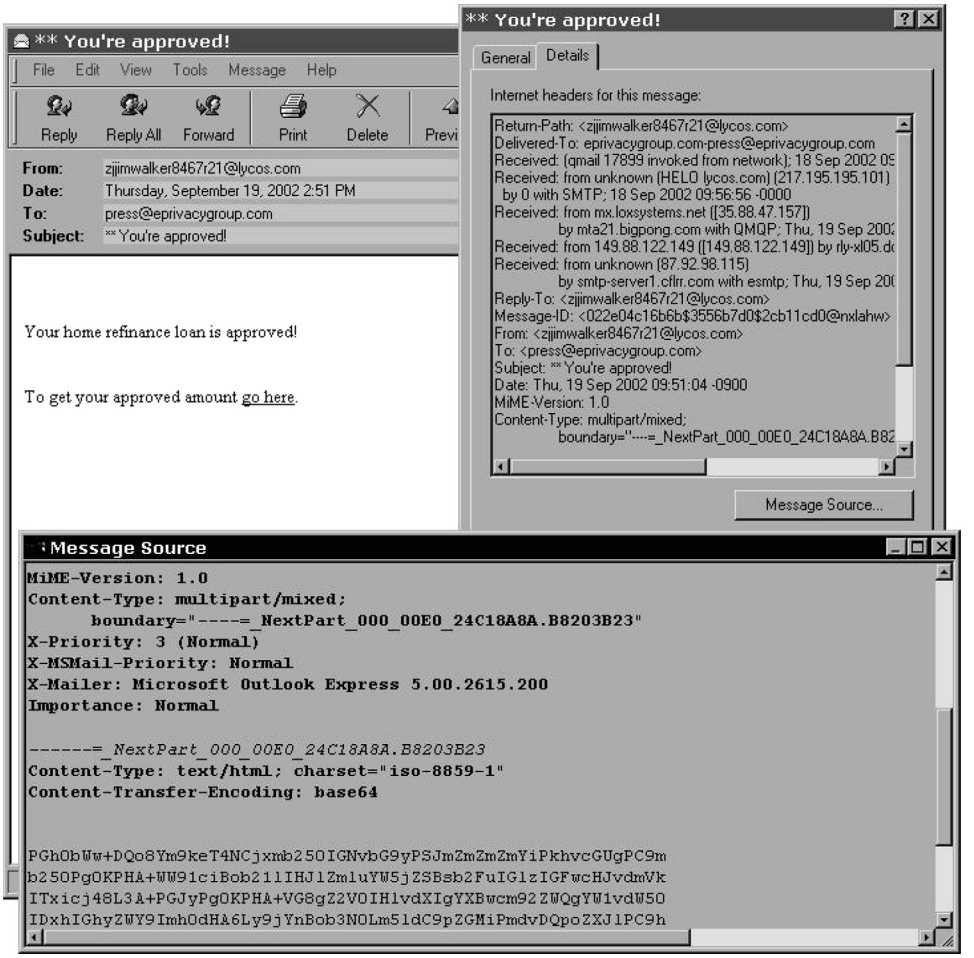


EXHIBIT 20.2 Viewing Message Header Details in Outlook Express

to press@eprivacygroup.com with the subject: *You're approved! As you may have guessed, press@eprivacygroup.com is not a real person. This is just an address that appears on a company Website as a contact for the press, and of course, nobody actually used this address on a mortgage application. The address was harvested by a program that automatically scours the Web for email addresses.

When you open this message in email clients (e.g., Outlook or Outlook Express) and use the File/Properties command or equivalent, you can click on the Details tab to see how the message made its way through the Internet. The first thing you see is a box labeled "headers," as shown in Exhibit 20.2. Reading this will tell you that the message was routed through several different email servers. Qmail is the company's mail program, which is the first instance of Received. The next three below that are intermediaries, until you get to the last one, smtp-server1.cflrr.com. That is an email server in Central Florida (cfl) on the Road Runner (rr) cable modem network, which supplies high-speed Internet access to tens of thousands of households.

So who sent this message? That is very hard to say. As you might expect, there is no such address as zjjimwalker8467r21@lycos.com. The best way to determine who sent a spam message like this is to examine the content. Spam cannot reel in suckers

unless there is some way for the suckers to contact the spammer. Sometimes this is a phone number, but in this message it is a hyperlink to a Web page. However, clicking links in spam messages is a dangerous way to surf—a much safer technique to learn more about links in spam is message source inspection. Outlook Express provides a Message Source view, but this may only show the encoded content and not be readable ASCII. Some email clients (e.g., Outlook) allow one to set a default of conversion of all messages into plain text; if one wishes to see images, one can have them downloaded on a per-message basis. In addition, many email clients can display the underlying details of any hyperlink when one hovers the cursor over the label for the link.

One way to get at the content of such messages is to forward them to a different email client, for example, Qualcomm's Eudora, then open the mailbox file with a text editor such as TextPad. This reveals the http reference for the link that this spammer wants the recipient to click. In this case, the sucker who clicks on the link that says "To get your approved amount go here" is presented with a form that is not about mortgages, but about data gathering. To what use the data thus gathered will be put is impossible to say, but a little additional checking using the *ping* and *whois* commands reveals that the Web server collecting the data is in Beijing. The chances of finding out who set it up are slim. The standard operating procedure is to set up and take down a Website like this very quickly, gathering as much data as you can before someone starts to investigate. However, over the last 10 years, the resources devoted to investigating spammers have been minimal compared to the resources that have gone into sending spam. Even when lawmakers can agree on what spam is and how to declare it illegal, government funds are rarely allocated for spam fighting. There have been a few high-profile arrests and prosecutions, often driven by large companies, such as Microsoft and AOL, but spamming remains a relatively low risk form of computer abuse.

If one can identify a link in spam, the Better-Whois service (http://betterwhois.com/) may be able to provide details of who is officially registered as owners and administrators of the site—assuming they have not hidden themselves behind anonymizing companies such as GoDaddy.com, which can conceal the true owners while providing an anonymized email address through which to attempt to communicate with the owners.

20.3 3 SPAM DEFINED.

The story of spam is a rags-to-riches saga of a mere nuisance that became a multibillion-dollar burden on the world's computing resources. For all the puns and jokes one can make about spam, it may be the largest single thief of computer and telecommunication resources since computers and telecommunications were invented.

Perhaps it is no surprise that spam has become a costly problem, because spam was the first large-scale computer threat to be purely profit-driven. The goal of most spam is to make money, and it cannot make money if people do not receive it. In fact, spam software will not send spam to a mail server that has a slow response time; it simply moves on to other targets that can receive email at the high message-per-minute rate needed for spam to generate income (based on 1 person in perhaps 100,000 actually responding to the spam message). If spam does not generate income, it becomes pointless because it takes money to send spam. You need access to servers and bandwidth (which you either have to pay for or steal). Ironically, a few simple changes to current email standards could put an end to most spam by more reliably identifying and authenticating email senders, a subject addressed later in this chapter.

20.3.1 Origins and Meaning of Spam (not SPAM®). SPAM® has been a trademark of Hormel Foods for over 70 years.[5] For reasons that will be discussed in a moment, the world has settled on *spam* as a term to describe unwanted commercial email. However, uppercase SPAM® is still a trademark, and associating SPAM®, or pictures of SPAM®, with something other than the Hormel product could be a serious violation of trademark law. Security professionals should take note of this. The word *Spam* is acceptable at the start of a sentence about unsolicited commercial email, but SPAM can be used for spam only if the rest of the text around it is uppercase, as in "TIRED OF SPAM CLUTTERING YOUR INBOX?"

The use of the word "spam" in the context of electronic messages is widely thought to stem from a comedy sketch in the twenty-fifth episode of the BBC television series *Monty Python's Flying Circus*.[6] First aired in 1968, before email was invented, the sketch featured a restaurant in which SPAM® dominated the menu. When a character called Mr. Bun asks the waitress, "Have you got anything without SPAM® in it?" she replies: "Well, there's SPAM®, egg, sausage and SPAM®, that's not got much SPAM® in it." The banter continues in this vein until Mr. Bun's exasperated wife screams, "I don't like SPAM®!" She is further exasperated by an incongruous group of Viking diners singing a song, the lyrics of which consist almost entirely of the word "SPAM®." A similar feeling of exasperation at having someone foist on you something that you do not want and did not ask for clearly helped to make "spam" a fitting term for the sort of unsolicited commercial email that can clutter inboxes and clog email servers.

In fact, the first use of spam as a term for abuse of networks did not involve email. The gory details of spam were carefully researched in 2003 by Brad Templeton, former chairman of the board of the Electronic Frontier Foundation.[7] According to Templeton, the origins lie in annoying and repetitive behavior observed in multiuser dungeons (also known by the acronym MUDs, an early term for a real-time, multiperson, shared environment). From there the term *spam* migrated to bulletin boards—where it was used to describe automated, repetitive message content—thence to USENET, where it was applied to messages that were posted to multiple newsgroups.

A history of USENET and its relationship to the history of spam is salutary for several reasons. First of all, spam pretty much killed USENET, which was once a great way to meet and communicate with other Internet users who shared common interests, whether those happened to be political humor or HTML coding. Newsgroups got clogged with spam to the point where users sought alternative channels of communication. In other words, a valuable and useful means of communication and an early form of social networking was forever tainted by the bad behavior of a few individuals who were prepared to flaunt the prevailing standards of conduct.

The second spam–USENET connection is that spam migrated from USENET to email through the *harvesting* of email addresses, a technique that spammers—the people who distribute spam—then applied to Web pages and other potential sources of target addresses. Spammers found that software could easily automate the process of reading thousands of newsgroup postings and extracting, or harvesting, any email addresses that appeared in them. It is important to remember that, although it might seem naive today, the practice of including one's email address in a message posted to a newsgroup was common until the mid-1990s. After all, if you posted a message looking for an answer, then including your email address made it easy for people to reply. It may be hard for some computer users to imagine a time when email addresses were so freely shared, but that time is worth remembering because it shows how easily the abuse of a system by a few can erode its value to the many. Harvesting, which resulted in receiving spam at an email address provided in a newsgroup posting for

the purposes of legitimate communication, was an enormous advance in abuse of the Internet and a notable harbinger of things to come. Elaborate countermeasures had to be developed, all of which impeded the free flow of communication. Much of the role fulfilled by newsgroups migrated to closed forums, where increasingly elaborate measures are used to prevent abuse.

20.3.2 Digging into Spam. Around 1996, as early adopters of email began experiencing increasing amounts of unsolicited email, efforts were made to define exactly what kind of message constituted spam. This was no academic exercise, and the stakes were surprisingly high. We have already employed one definition: unsolicited commercial email (sometimes referred to as UCE). This definition seems simple enough, capturing as it does two essential points: Spam is email that people did not ask to receive, and spam is somehow commercial in nature.

However, while UCE eventually became the most widely used definition of spam, it did not satisfy everyone. Critics point out that it does not address unsolicited messages of a political or noncommercial nature (such as a politician seeking your vote or a charity seeking donations). Furthermore, the term "unsolicited" is open to a lot of interpretation, a fact exploited by early mass emailings by mainstream companies whose marketing departments used the slimmest of pretexts to justify sending people email. As the universe of Internet users expanded in the late 1990s, encompassing more and more office workers and consumers, three things about spam became increasingly obvious:

1. People did not like spam.
2. Spammers could make money.
3. Legitimate companies were tempted to send unsolicited email.

20.3.2.1 Get-Rich-Quick Factor. How much money could a spammer make? Consider the business plan of one Arizona company, C.P. Direct, that proved very profitable until it was shut down in 2002 by the U.S. Customs Service and the Arizona Department of Public Safety. Here are some of the assets that authorities seized:

- Nearly $3 million in cash plus a large amount of expensive jewelry
- More than $20 million in bank accounts
- Twelve luxury imported automobiles (including eight Mercedes, plus assorted models from Lamborghini, Rolls-Royce, Ferrari, and Bentley)
- One office building and assorted luxury real estate in Paradise Valley and Scottsdale

These profits were derived from the sale of $74 million worth of pills that promised to increase the dimensions of various parts of the male and female anatomy. The company had used spam to sell these products, but in fact it was not shuttered for sending spam, the legal status of which remains ambiguous to this day, having been defined differently by different jurisdictions (not least of which was the U.S. Congress, which arguably made a hash of it in 2004).

C.P. Direct crossed several lines, not least of which was the making of false promises about its products (none of which it ever tested and all of which turned out to contain the same ingredients, regardless of which part of the human anatomy they were promised

to enhance). The company compounded its problems by refusing to issue refunds when the products did not work as claimed. However, rather than discourage spammers, this case proved that spam can make you rich, quick. Indeed, the hope of getting rich quick remains the main driver of spam. The fact that a handful of people were prosecuted for conducting a dubious enterprise by means of spam was not perceived as a serious deterrent.

20.3.2.2 Crime and Punishment. Spammers keep on spamming because risks are perceived to be small relative to both the potential gains and the other options for getting rich quick. The two people at the center of the C.P. Direct case, Michael Consoli and his nephew and partner, Vincent Passafiume, admitted their guilt in plea agreements signed in August 2003; but they were out of jail before May 2004, and they seemed to have suffered very little in the shame department. Two years after their release, the pair asked the state Court of Appeals to overturn the convictions and give back whatever was left of the seized assets.

Consider what has happened with Jeremy Jaynes, named as one of the world's top 10 spammers in 2003 by Spamhaus, a spam-fighting organization discussed later in Section 20.4.1 of this chapter. When he was prosecuted for spamming, Jaynes was thought to be sending out 10 million emails a day. How much did he earn from this? Prosecutors claimed it was about $750,000 a month. In 2004, Jaynes was convicted of sending unsolicited emails with forged headers, in violation of Virginia law. He was sentenced to nine years in prison. However, in September 2008, Jaynes, who had not served any time and remained free on bail (set at less than two months' worth of his spam earnings) had his conviction overturned by the Virginia Supreme Court, which ruled that the state antispam law was unconstitutional[8]—a ruling attacked as ill-founded by some critics specializing in antispam legislation.[9]

Another notorious spammer was Sanford "Spamford" Wallace, founder of Cyber Promotions in the 1990s, which actively used spam as a commercial service. In October 2009, Facebook won a civil suit against him for sending fraudulent messages to its users.[10] Wallace was ordered to pay a $711 million fine—which he was unlikely to pay because he filed for bankruptcy in June 2009.

In August 2011, he was indicted in the federal court in San Jose on

> … multiple counts of fraud and related activity in connection with electronic mail. Wallace was also charged with three counts of intentional damage to a protected computer and two counts of criminal contempt. According to the indictment, from approximately November 2008 through March 2009, Wallace executed a scheme to send spam messages to Facebook users. Those messages compromised approximately 500,000 legitimate Facebook accounts, and resulted in more than 27 million spam messages being sent through Facebook's servers. The indictment alleges that Wallace sent spam messages to Facebook users during three time periods: First, on or about Nov. 5, 2008, and continuing until approximately Nov. 6, 2008, Wallace accessed Facebook's computer network to initiate the transmission of a program that resulted in more than 125,000 spam messages being sent to Facebook users; Second, on Dec. 28, 2008, Wallace accessed Facebook's computer network to initiate the transmission of a program that resulted in nearly 300,000 spam messages being sent to Facebook users; Third, on Feb. 17, 2009, Wallace accessed Facebook's computer network to initiate the transmission of a program that resulted in more than 125,000 spam messages being sent to Facebook users.[11]

At the time of writing (May 2013), USA v. Sanford Wallace was scheduled to open on Monday June 3, 2013, in the court of Judge Edward J. Davila in the San Jose Courthouse.[12]

20.3.2.3 A Wasteful Game. In many ways, spam is the heir to the classic sucker's game played out in classified advertisements that promise to teach you how to "get cash in your mailbox." The trick is to get people to send you money to learn how to get cash in their mailbox. If a spammer sends out 25 million email messages touting a product, she *may* sell enough products to make a profit, but there are real production costs associated with a real product. In contrast, if she sends out enough messages touting a list of 25 million proven email addresses for $79.95, she may reel in enough suckers willing to buy that list and make a significant profit, because the list costs essentially nothing to generate. The fact that most addresses on such lists turn out to be useless does not seem to stop people from buying or selling them.

One form of spam that does not rely on selling a product is the *pump-and-dump* stock scam. Sending out millions of messages telling people how hot an obscure stock is about to become can create a self-fulfilling prophecy. Here is how it works:

- Buy a lot of shares on margin or a lot of shares in a company that is trading at a few pennies a share.
- Spam millions of people with a message talking up the stock you just bought.
- Wait for the price of the shares to go up, and then sell your shares for a lot more than you paid for them.

Such a scheme breaks a variety of laws but can prove profitable if you are not caught (you can hide your identity as a sender of spam); nevertheless, regulators still can review large buy and sell orders for shares referred to in stock spams.

The underlying reasons for the continuing rise in spam volume can be found in the economics of the medium. Sending out millions of email messages costs the sender very little. An ordinary personal computer (PC) connected to the Internet via a $10-per-month dial-up modem connection can pump out hundreds of thousands of messages a day; a small network of PCs connected via a $50-per-month cable modem or digital subscriber line (DSL) can churn out millions. Obviously, the economic barrier to entry into get-rich-quick spam schemes is very low. The risk of running into trouble with the authorities is also very low. The costs of spam are borne downstream, in a number of ways, by several unwilling accomplices:

Email Recipient

- Spends time separating junk email from legitimate email. Unlike snail mail, which typically is delivered and sorted once per day, email arrives throughout the day and night. Every time you check it, you face the time-wasting distraction of having to sort out spam.
- Pays to receive email. There are no free Internet connections. When you connect to the Internet, somebody pays. The typical consumer at home pays a flat rate every month, but the level of that rate is determined, in part, by the volume of data that the Internet service provider (ISP) handles, and spam inflates that volume, thus inflating cost.

Enterprise

- Loses productivity because employees, many of whom must check email for business purposes, are spending time weeding spam out of their company email

inbox. Companies that allow employees to access personal email at work also pay for time lost to personal spam weeding.

- Wastes resources because spam inflates bandwidth consumption, processing cycles, and storage space.

Internet Service Providers (ISPs) and Email Service Providers (ESPs)[13]

- Wastes resources on handling spam that inflates bandwidth consumption, processing cycles, and storage space.
- Have to spend money on spam filtering, block list administration, spam-related customer complaints, and filter/block-related complaints.
- Have to devote resources to policing their users to avoid getting block-listed. (There is more about filters and block lists in the next section.)

Two other economic factors are at work in the rise of spam: hard times and delivery rates. When times are tough, more people are willing to believe that get-rich-schemes like spamming are worth trying, so there are more spammers (tough times affect the receiving end as well, with recipients more willing to believe fraudulent promises of lotteries won and easy money to be made). When delivery rates for spam go down—due to the use of spam filters and other techniques that are discussed in Section 20.4.5—spammers compensate by sending even more spam.

20.3.2.4 How Big Is the Spam Problem? There is some controversy about the proportion of email classified as spam. By 2006, spam was said to be consuming over 90 percent of email resources worldwide.[14] This was a staggering level of system abuse by any standard. However, when a handful of security professionals had claimed, a decade earlier, that spam was a computer-security threat, they were met with considerable skepticism and some suspicion, perhaps in part because of the anticlimax of Y2K; also, there was doubtless an element of the recurrent suspicion that security professionals trumpet new threats to drum up business—a strange notion, given the perennial abundance of opportunities for experts in a field that persistently reports near-zero levels of unemployment.

A useful historical perspective on spam's depressing impact on email is provided by the reports freely available at the MessageLabs Website. Its annual reports provide conservative estimates of the growth in spam; for example, the 2007 annual report[15] showed total spam hovering around 85 percent of all emails from 2005 through 2007, with new varieties (those previously unidentified by type or source) keeping fairly steady at around 75 percent of all email.

Symantec reported about spam as over 90 percent of total email traffic in 2009[16] and 2010,[17] but the proportion began dropping over the next few years. In June 2011, Symantec reported a spam rate of about 73 percent of total email[18]; by December 2011, they reported a further drop to about 70 percent.[19] According to Kaspersky Lab in early 2013, "… the share of spam in email traffic decreased steadily throughout 2012 to hit a five-year low. The average for the year stood at 72.1%—8.2 percentage points less than in 2011. Such a prolonged and substantial decrease in spam levels is unprecedented."[20] A different study in January 2013 suggested that only about 60 percent of all email was spam in 2012.[21] Estimates by other experts suggested that only about 15 percent of the total spam was getting through all the spam filters at ISP and application levels.[22]

Although most companies and consumers probably will concur that spam has grown from a mere annoyance to a huge burden, some people will say spam is not a big problem. These include:

- Consumers who have not been using email for very long
- Office users who do not see the spam addressed to them due to some form of antispam device or service deployed by their company

These perceptions obscure the fact, noted earlier, that spam consumes vast amounts of resources that could be put to better use. The consumer might get cheaper, better Internet service if spam was not consuming so much bandwidth, server capacity, storage, and manpower. In the author's experience working with a regional ISP, growth in spam volume is directly reflected in server costs. The ISP had to keep adding servers to handle email. By the time it got to four servers, 75 percent of all the email was determined to be spam. The ISP had incurred server costs four times greater than needed to handle legitimate email. Furthermore, even with four servers, spikes in spam volumes were causing servers to crash, incurring the added cost of service calls in the middle of the night, not to mention the loss of subscribers annoyed by outages.

Spam has a direct effect on infrastructure spending. Storage is one of the biggest hardware and maintenance costs incurred by email. If even 70 percent of all email is unsolicited junk, then companies that process email are spending much more on storage than they would if all email was legitimate. The productivity hit for companies whose employees waste time weeding spam out of their inboxes is also enormous. Beyond these costs, consider the possibility that spam actually acts as a brake on Internet growth rates, having a negative effective on economies, such as America's, that derive considerable strength from Internet-related goods and services.

Although the Internet appears to grow at a healthy clip, it may be a case of doing so well that it is hard to know how badly we are doing. Perhaps the only reason such an effect has not yet been felt is that spam's impact on new users is limited. As noted earlier, when new users first get email addresses, they typically do not get a lot of spam. According to some studies, it can take six to twelve months for spammers to find an email address, but when they do, the volume of spam to that address can increase very quickly. This leads some people to cut back on their use of email and the Internet.

In "The Economics of Spam," published in the Summer 2012 issue of the *Journal of Economic Perspectives*, Justin M. Rao (Microsoft) and David H. Reiley (Google), both formerly employees of Yahoo! Research, discuss the *externality* of spam—the use of victims' resources to support profit for the criminals. They write,

> We estimate that American firms and consumers experience costs of almost $20 billion annually due to spam. Our figure is more conservative than the $50 billion figure often cited by other authors, and we also note that the figure would be much higher if it were not for private investment in anti-spam technology by firms… . On the private-benefit side, based on the work of crafty computer scientists who have infiltrated and monitored spammers' activity… we estimate that spammers and spam-advertised merchants collect gross worldwide revenues on the order of $200 million per year. Thus, the "externality ratio" of external costs to internal benefits for spam is around 100:1.[23]

20.3.3 Spam's Two-Sided Threat. When mail servers slow down, falter, and finally crash under an onslaught of spam, the results include lost messages, service interruptions, and unanticipated helpdesk and tech support costs. Sales are missed.

Customers do not get the service they expect. The cost of keeping spam out of your inbox and your enterprise is one thing, the cost of preventing spam from impacting system availability and business operations is another; but there is another side of the spam threat, the temptation to become an abusive mass mailer, otherwise known as a spammer. This is something that spam has in common with competitive intelligence, otherwise known as industrial espionage. When done badly, mass mailings, like competitive intelligence, can tarnish a company's reputation.

Leaving aside the matter of the increasingly nasty payloads delivered with spam, things like Trojans and phishing attacks, even plain old male enhancement spam constitutes a threat to both network infrastructure and productivity. For a start, spam constitutes a theft of network resources. The inflationary effect on server budgets has already been mentioned. The negative impact on bandwidth may be less obvious but is definitely real. In 2002 and 2003, the author was involved in the beta testing of a prototype network-level antispam router. It was not unusual for a company installing this device to discover that spam had been consuming from two-thirds to three-quarters of its network bandwidth. Whether this impact was seen as performance degradation or cost inflation, very few companies were willing to remove this device once it had been installed. In other words, when companies see what their network performance and bandwidth cost is like when spam is taken out of the equation, they realize just what a negative impact spam has. This is something that might otherwise be hard to detect given that spam has risen in volume over time.

An even more dramatic illustration of the damage that spam can cause comes when a network is targeted by a really big spam cannon (a purpose-built configuration of MTA devices connected to a really big broadband connection; for example, a six-pack of optimized MTAs can fire off 3.6 million messages an hour). The effect can be to crash the receiving mail servers, with all of the attendant cost and risk that involves. One way to prevent this happening, besides deployment of something like an antispammer router, is to sign up for an antispam service, which intercepts all of your incoming email and screens out the spam. Such a solution addresses a number of email-related problems, but at considerable ongoing cost, which still constitutes a theft of resources by spammers.

The company Commtouch.com provides a spam cost calculator that produces some interesting numbers.[24] Consider these input values for a medium-size business:

Employees: 800

- Average annual salary: $45,000
- Average number of daily emails per recipient: 75
- Average percentage of email that is spam: 80 percent

According to the calculator, total annual cost of spam to this organization, which is assumed to be deploying no antispam measures, is just over $1 million. This is based on certain assumptions, such as the time taken to delete spam messages, but in general, it seems fairly realistic. Of course, the size of the productivity hit caused by the spam that makes it to an employee inbox has been hotly debated over the years, but it is clearly more than negligible and not the only hit. Even if antispam filtering is introduced, there will still be a need to review or adjust the decisions made by the filter to ensure that no important legitimate messages are erroneously quarantined or deleted. In other words, even if the company spends $50,000 per year on antispam filtering, it will not reclaim all of the $1 million wasted by spam.

20.3.3.1 Threat of Outbound Spam. Even as organizations like the Coalition Against Unsolicited Commercial Email (CAUCE) were trying to persuade companies that spamming was an ill-advised marketing technique that could backfire in the form of very annoyed recipients, and just as various government entities were trying to create rules outlawing spam, some companies were happy to stretch the rules and deluge consumers with email offers, whether those consumers had asked for them or not. This led some antispammers to condemn all companies in the same breath. The author's own experience, working with large companies that have respected brand names, was that none of them actually wanted to offend consumers. Big companies always struggle to rein in maverick marketing activities, and mass emailing is very tempting when a rogue employee or simply an uninformed one are under pressure to produce sales; but upper management is unlikely to condone anything that could be mistaken for spamming.

The motives for corporate responsibility in email are not purely altruistic. Smart companies can see that the perpetuation of disreputable email tactics only dilutes the tremendous potential of email as a business tool. Whatever one thinks of spam, there is no denying that, as a business tool, bulk email is powerful. It is also seductive. When one has a story to tell or a product to sell, and a big list of email addresses is just sitting there, bulk email can be very tempting. Naïve users can find themselves thinking "Where's the harm?" and "Who's going to object?" But unless they have documented permission to send their message to the people on that list, the smart business decision is to resist the temptation. Remember, it only takes one really ticked-off recipient of your unsolicited message to ruin your day.

20.3.3.2 Mass Email Precautions. One of the most basic business email precautions is this: Never send a message unless you are sure you know what it will look like to the person who receives it. This covers the formatting, the language you use, and, above all, the addressing. If you want to address the same message to more than one person at a time, you have three main options, each of which should be handled carefully:

1. Place the email addresses of all recipients in the To field or the Copy (Cc) field so all the recipients will be able to see the addresses of the other people to whom you sent the message. This is sometimes appropriate for communications within a small group of people.
2. If the number of people in the group exceeds about 20, or if you do not want everyone to know who is getting the message, move all but one of them to the Blind Copies (Bcc) field. The one address in the To field may be your own. If the disclosure of recipients is likely to cause any embarrassment whatsoever, do a test mailing first. Send a copy of the message to yourself and to at least one colleague outside the company, and then have the message looked at to make sure the Bcc entries were made correctly.
3. To handle large groups of recipients, or to personalize one message to many recipients, use a specialized application like Group Mail that can reliably build individual, customized messages to each person on a list. This neatly sidesteps errors related to the To and Cc fields. Group Mail stores email addresses in a database and builds messages on the fly using a merge feature like a word processor. You can insert database fields in the message. For example, a message can refer to the recipient by name. The program also offers extensive testing

of messages, so that you can see what recipients will see before you send any messages. And the program has the ability to send messages in small groups, spread over time, according to the capabilities of your Internet connection.

Whether you use a program like Group Mail for your legitimate bulk email, or something even more powerful, depends on several factors, such as the size of your organization, the number of messages you need to send, and your privacy policy. Privacy comes into play because some software used to send email, such as Group Mail, allows the user of the program access to the database containing the addresses to which the mail is being sent. This is generally not a problem in smaller companies, or when the database consists solely of names and addresses without any special context; but it can be an issue when the database contains sensitive information or the context is sensitive. For example, a list of names and addresses can be sensitive if the person handling the list knows, or can infer, that they belong to patients undergoing a certain kind of medical treatment.

You may not want to allow system operators or even programmers to have access to sensitive data simply because they are charged with sending or programming messages. Fortunately, mailing programs can be written that allow an operator to compose mail and send it without seeing the names and addresses of the people to whom it is being sent. Test data can, and should, be used to proof the mailing before it is executed.

Another basic email precaution is never to send any message that might offend any of the recipients. In choosing your wording, your design, your message, know your audience. Use particular caution when it comes to humor, politics, religion, sex, or any other sensitive subject. When using email for business email, it is better to stand accused of a lack of humor rather than a lack of judgment.

Respect people's preferences, if you know them, for content. If people have expressed a preference for text-only messages, do not send them HTML and hope they decide to change their minds. Ask first because the forgiveness-later path is not cost effective when you have to deal with thousands of unhappy recipients calling your switchboard. When you do use HTML content, still try to keep size to a minimum, unless you have recipients who specifically requested large, media-rich messages.

Antispam measures deployed by many ISPs, companies, and consumers sometimes produce false positives, flagging legitimate email as spam, potentially preventing your email from reaching the intended recipients, even when they have asked to receive your email. If your company's email is deemed to be particularly egregious spam, the server through which it is sent is likely to be blocked. If this is your company's general email server, blocking could affect delivery of a lot more than just the spam. If the servers through which your large mailing is sent belong to a service provider, and its servers get blocked, that could be a problem for you too. And if you use a service provider to execute the mailing for you but fail to choose wisely, your mail may be branded as spam just because of the bad reputation of the servers through which it passes. Be aware that using spam to advertise your products could place you in a violation of the CAN-SPAM Act of 2003, even if you do not send the spam yourself. (See Section 20.4.9 for more on CAN-SPAM.)

What can be done to prevent messages that are not spam from falling victim to antispam measures? Adherence to responsible email practices is a good first step. Responsible management of the company's email servers will also help, as will selection of reputable service providers. You should also consider tasking someone with tracking antispam measures, to make sure that your email is designed to avoid, as much as

possible, any elements of content or presentation that are currently being flagged as spam.

To make sure your company is associated with responsible email practices, become familiar with the "Six Resolutions for Responsible E-Mailers." These were created by the Council for Responsible Email (CRE), which was formed under the aegis of the Association for Interactive Marketing (AIM), a subsidiary of the Direct Market Association (DMA). Some of the country's largest companies, and largest legitimate users of email, belong to these organizations, and they have a vested interest in making sure that email is not abused. Here are the six resolutions:

1. Marketers must not falsify the sender's domain name or use a nonresponsive IP address without implied permission from the recipient or transferred permission from the marketer.
2. Marketers must not purposely falsify the content of the subject line or mislead readers from the content of the email message.
3. All bulk email marketing messages must include an option for the recipient to unsubscribe (be removed from the list) from receiving future messages from that sender, list owner, or list manager.
4. Marketers must inform the respondent at the time of online collection of the email address for what marketing purpose the respondent's email address will be used.
5. Marketers must not harvest email addresses with the intent to send bulk unsolicited commercial email without consumers' knowledge or consent. (Harvest is defined as compiling or stealing email addresses through anonymous collection procedures such as via a Web spider, through chat rooms, or other publicly displayed areas listing personal or business email addresses.)
6. The CRE opposes sending bulk unsolicited commercial email to an email address without a prior business or personal relationship. (Business or personal relationship is defined as any previous recipient-initiated correspondence, transaction activity, customer service activity, third-party permission use, or proven off-line contact.)

These six resolutions may be considered a reasonable middle ground between antispam extremists and marketing-at-all-costs executives. If everyone abided by these resolutions, there would be no spam, at least according to most people's definition of spam. After all, if nobody received more than one or two messages per week that were unwanted and irrelevant, antispam sentiment would cool significantly.

20.3.3.3 Appending and Permission Issues. Some privacy advocates have objected to the sixth resolution because it permits email *appending*. This is the practice of finding an email address for a customer who has not yet provided one. Companies such as Yesmail and AcquireNow will do this for a fee. For example, if you are a bank, you probably have physical addresses for all your customers, but you may not have email addresses for all of them. You can hire a firm to find email addresses for customers who have not provided one. However, these customers may not have given explicit permission for the bank to contact them via email, so some people would say that sending email messages to them is spamming.

Whether you agree with that assessment or not, several factors need to be assessed carefully if your company is considering using an email append service. First of all,

make sure that nothing in your privacy policy forbids it. Next, think hard about the possible reaction from customers, bearing in mind that email appending is not a perfect science. For more on how appending works, enter "email append" as a search term in Google—you will find a lot of companies offering to explain how they do it, matching data pulled from many different sources using complex algorithms.

Another concern that must be considered is that some messages will go to people who are not customers. For that reason, you probably want to make your first contact a tentative one, such as a polite request to make further contact. Then you can formulate the responses to build a genuine opt-in list. As you might expect, you can outsource this entire process to the append service, which will have its own, often automated, methods of dealing with bounced messages, complaints, and so on.

Do not include any sensitive personal information in the initial contact, since you have no guarantee that bob.jones@majorfreemail.net is the Robert Jones you have on your customer list. When Citibank did an append mailing in the summer of 2002, encouraging existing customers to use the bank's online account services, it came in for some serious criticism. Although the bank was not providing immediate online access to appended customers, the mere perception of this, combined with a number of mistaken email identities, produced negative publicity. Here is what Citibank said in the message it sent to people it believed to be customers, even though these people had not provided their email address directly to the bank:

> Citibank would like to send you email updates to keep you informed about your Citi Card, as well as special services and benefits… With the help of an email service provider, we have located an email address that we believe belongs to you.

Although this message is certainly polite, it clearly raised questions in the minds of recipients. Two questions could undermine appending as a business practice: Where exactly is this service provider looking for these email addresses? Why doesn't the company that wants my email address just write and ask for it? The fact is, conversion rates from email contact are higher than from snail mail, so the argument for append services is that companies that use them move more quickly to the cheaper and better medium of email than those that do not. The counterpoint is that too many people will be offended in the process.

Consider to what lengths you are stretching the prior business relationship principle cited in the sixth responsible email resolution. A bank probably has a stronger case for appending email addresses to its account holder list than does a mail order company that wants to append a list of people who requested last year's catalog. The extent to which privacy advocates accept or decry the concept of prior business relationship is largely dependent on how reasonable companies are in their interpretation of it.

Marketing to addresses that were not supplied with a clear understanding that they would be used for such purposes is not advisable. Depending on your privacy statement, it could be a violation of your company's privacy policy. Going ahead with such a violation could not only annoy customers, but also draw the attention of industry regulators, such as the Federal Trade Commission (FTC). Of course, you will have to decide for yourself if it is your job or responsibility to point this out to management. And be sure to provide a simple way for recipients to opt out of any further mailings. (A link to a Web form is best for this. Avoid asking the recipient to reply to the message. If the email address to which you mailed the message is no longer their primary address, they may have trouble opting out.)

20.4 FIGHTING SPAM. The fight against incoming spam can be addressed both specifically, as in "What can I do to protect my systems from spam?", and in general, "What can be done to prevent spam in general?" Of course, if the activity of spamming were to be eliminated, everyone would have one less threat to worry about.

20.4.1 Enter the Spam Fighters. In the five-year period from 1997 to 2002, a diverse grouping of interests and organizations fought to make the world aware of the problem of spam, and to encourage counter-measures. The efforts of CAUCE, which continue to this day, spurred lawmakers into passing antispam legislation and helped bring order to the blacklisting process by which spam-sending servers are identified. In 1998, the *Spamhaus Project*, a volunteer effort founded by Steve Linford, began tracking spammers and spam-related activity. (The name comes from a pseudo-German expression, coined by Linford, to describe any ISP or other firm that sends spam or willingly provides services to spammers.)

The Anti-Phishing Working Group (APWG) was founded in 2003 and is one of the most active and productive anti-phishing organizations today:

> The APWG is a worldwide coalition unifying the global response to cybercrime across industry, government and law-enforcement sectors. APWG's membership of more than 2000 institutions worldwide is as global as its outlook, with its directors, managers and research fellows advising: national governments; global governance bodies like ICANN; hemispheric and global trade groups; and multilateral treaty organizations such as the European Commission, Council of Europe's Convention on Cybercrime, United Nations Office of Drugs and Crime, Organization for Security and Cooperation in Europe and the Organization of American States. Membership is open to financial institutions, retailers, solutions providers, ISPs, telcos, defense contractors, law enforcement agencies, trade groups, treaty organizations and government agencies. APWG member benefits include: clearinghouses of cybercrime event data; cybercrime response utilities for professionals from the private, public, and NGO sectors who combat cybercrime; community-building conferences for cybercrime management professionals; public education utilities for cybercrime prevention; standards development for cybercrime data exchange and programs for promotion of cybercrime research.[25]

Commercial antispam products started to appear in the late 1990s, starting with spam filters that could be used by individuals. Filtering supplied as a service to enterprises appeared in the form of Brightmail, founded in 1998 and now owned by Symantec, and Postini, founded in 1999 and now owned by Google. A host of other solutions, some free and open source, others commercial, attempted to tackle spam from several different directions. Nevertheless, despite the concerted efforts of volunteers, lawmakers, and entrepreneurs, spam has continued to sap network resources while evolving into part of a subculture of system abuse that includes delivering payloads that do far worse things than spark moral outrage.

20.4.2 A Good Reputation? Since spam is created by humans, it is notoriously difficult, if not impossible, for computers to identify spam with 100 percent reliability. This fact led some spam fighters to consider an alternative approach: reliably identifying legitimate email. If an email recipient or a Mail Transfer Agent could verify that certain incoming messages originated from legitimate sources, all other email could be ignored. One form of this approach is the challenge-response system, perhaps the largest deployment being the one that Earthlink rolled out in 2003. When you send a message to someone at Earthlink who is using this system, and who has not received a message from you before, your message is not immediately delivered. Instead, Earthlink sends an email asking you to confirm your identity, in a way that

would be difficult for a machine to fake. The recipient is then informed of your message and decides whether to accept it or not. Unfortunately, this approach can be problematic if the user does a lot of e-commerce involving email coming from many sources that are essentially automated responders (which cannot pass the challenge). Manually building the whitelist of allowed respondents can be tiresome, and failure to do so may mean some messages do not get through (e.g., if the sender is a machine that does not know how to handle the challenge). One solution to this problem is to compile an independent whitelist of legitimate emailers whose messages are allowed through without question. This is the reputational approach to spam fighting, and it works like this:

- Bank of America pledges never to spam its customers and send them only email they sign up for (be it online statement notification or occasional news of new bank services).
- Bank of America email is always fast-tracked through ISPs and not blocked as spam.
- Bank of America remains true to its pledge because its reputation as a legitimate mailer enables it to conduct email activities with greater efficiency.

There are considerable financial and logistical obstacles to making such a system work, and it tends to work better with bigger mailers. Adoption also faced skepticism from some privacy and antispam advocates who suspected that the goal of such systems was merely to legitimize mass mailings by companies that still did not understand the need to conduct permission-only mailings. The relentless assault of spam on the world's email systems eventually may lead all legitimate companies to foreswear unsolicited email and thus make a reputational system universally feasible. Tough economic times, though, may tempt formerly clean companies to turn to spam in a desperate effort to boost flagging sales.

Another and more universal method of excluding spam would be for ISPs to deliver, and consumers to accept, only those emails that were stamped with a verifiable cryptographic seal of some kind. A relatively simple automated system to accomplish this was developed in 2001 by a company called ePrivacy Group, and it proved very successful in real-world trials conducted by MSN and several other companies. If adopted universally, such a system could render spamming obsolete, but that goal proved to be unattainable. The success of this approach depended on widespread deployment, and the project ultimately was doomed by infighting between the larger ISPs, despite ePrivacy Group's willingness to release the underlying technology to the public domain.

The ultimate in reputation-based approaches to solving the spam problem may well be something that information security experts have urged for years: widespread use of email encryption. If something like Secure/Multipurpose Internal Mail Extensions (S/MIME) was universally implemented, everyone could ignore messages that were not signed by people from whom they were happy to receive email. Of course, the fact that email encryption can be made to work reliably is not proof that it always will, and the encryption approach does not, in itself, solve the fundamental barrier to any universal effort at outlawing spam: the willingness of everyone to participate. One of the qualities that led email to become the most widely used application on the Internet, the lack of central control, is a weakness when it comes to effecting any major change to the way it operates.

20.4.3 Relaying Trouble When considering the problem of spam in general, the question is why ISPs allow people to send spam. The fact is, many do not. Spamming is a violation of just about every ISP's terms of service. Accounts used for spamming are frequently closed. However, there are some exceptions. As you might imagine, a few ISPs allow spam as a way to get business (albeit business that few others want). Some of these ISPs use facilities in fringe countries to avoid regulatory oversight. Furthermore, some spammers find it easier, and cheaper, to steal service and use unauthorized access to other people's servers to send their messages. A phenomenon of email is mail relaying. According to Janusz Lukasiak of the University of Manchester, mail relaying occurs "when a mail server processes a mail message from an unauthorized external source with neither the sender nor the recipient being a local user. The mail server is an entirely unrelated third party to this transaction and the message should not be passed via the server."[26] The problem with unauthenticated third parties is that they can hide their identity.

In the early days of the Internet, many servers were left *open* so that people could conveniently relay their email through them at any time. It was quite acceptable for individuals to send their email through just about any mail server, since the impact on resources was minimal. Abuses by spammers, whose mass emailings *do* impact resources, led to ISPs instituting restrictions. Some ISPs require the use of port 587 for SMTP authentication. Others require logging in to the POP server to collect incoming mail before sending any out. Since that requires a user name and password, it prevents strangers from sending email through the server.

Open relays are now frowned on. However, relaying continues to occur, partly because configuring servers to prevent it requires effort. Also, spammers keep finding new ways to exploit resources that are not totally protected. For example, port 25 filtering, instituted to reduce spamming, can be bypassed with tricks like asynchronous routing and proxies. A relatively recent development is the use of botnets, groups of compromised computers, to deliver spam. (For more about botnets see Chapters 15, 16, 17, 30, 32, and 41 in this *Handbook.*)

A war is continually being waged between spammers and ISPs, and that war extends to the ISPs' ISP—most Internet service providers actually get their service from an even larger service provider, companies like AT&T, MCI, and Sprint, which are the backbone of the Internet. How much trouble do these companies have with spam? As far back as 2002, network abuse (spam) generated 350,000 trouble tickets each month at a single carrier. These companies are working hard to prevent things like mail relaying and to defeat the latest tricks that spammers have devised to get around their preventive measures.

20.4.4 Black Holes and Block Lists Black hole lists, or block lists, catalog server IP addresses from ISPs whose customers are deemed responsible for spam and from ISPs whose servers are hijacked for spam relaying. ISPs and organizations subscribe to these lists to find out which sending IP addresses should be blocked. The receiving end, such as the consumer recipient's ISP, checks the list for the connecting IP address. If the IP address matches one on the list, then the connection gets dropped before accepting any traffic. Other ISPs choose simply to ignore, or black hole, IP packets at their routers. Among the better known block lists are RBL, otherwise known as MAPS Realtime Blackhole List, Spamcop, and Spamhaus.[27]

How does an entity's IP address get on these lists? If an ISP openly permits spam, or does not adequately protect its resources against abuse by spammers, it will likely be reported to the list by one or more recipients of such spam. Reports are filed by people

who take the time to examine the spam's header, identify the culpable ISP, and make a nomination to a block list. Different block lists have different standards for verifying nominations. Some test the nominated server; others take into account the number of nominations. If an ISP, organization, or individual operating a mail server finds itself on the list by mistake, it can request to be removed, which usually involves a test by the organization operating the block list.

Note that none of this block-list policing of spam is official. All block lists are self-appointed and self-regulated. They set, and enforce, their own standards. The only recourse for entities that feel they have been unfairly blocked—and there have been plenty of these over the years—is legal action. Some block lists are operated outside of the United States, but if an overseas organization block-lists a server located in the United States, it probably can be sued in a U.S. court. However, it is important to note that the blocking is not done by the operator of the block list, it is done by ISPs that subscribe to, and are guided by, the lists.

20.4.5 Spam Filters Block-list systems filter out messages from certain domains or IP addresses, or a range of IP addresses; they do not examine the content of messages. Filtering out spam based on content can be done at several levels.

20.4.5.1 End User Filters Spam filtering probably began at the client level, and even today many email users perform manual negative filtering for spam, identifying spam by a process of elimination. This is easy to do with any email application that allows user-defined filters or rules to direct messages to different inboxes or folders. Many people have separate mailboxes into which they filter all of the messages they get from their usual correspondents, friends, family, colleagues, subscribed newsletters, and so on. This means that whatever is left in the in basket is likely to be spam, with the notable exception of messages from new correspondents for whom a separate mailbox and filter has not yet been created.

To perform a filter that positively identifies spam, elements common to spam messages need to be identified. Many products do this, and they usually come with a default set of filters that look for things like From addresses that contain a lot of numbers and Subject text that contains a lot of punctuation characters—spammers often add these in an attempt to defeat filters based on specific text, so the Subject line "You're Approved" might be randomly concatenated with special characters and spaces like this: "**You~re Ap proved!**"

In fact, the tricks and tweaks used by spammers in crafting messages designed to bypass filters are practically endless. Nevertheless, even the spam filtering built into some basic email programs offers a useful line of defense. For example, in Exhibit 20.3 you can see the control panel for spam filtering in the Eudora email application, which places mail in a Junk folder based on a score derived from common spam indicators, giving users control over how high they want to set the bar.

Unfortunately, some of the language found in spam also appears in legitimate messages, such as "You are receiving this message because you subscribed to this list" or "To unsubscribe, click here." This means that the default setting in a spam filter is likely to block some messages that you want to get. The answer is to either weaken the filtering or create a whitelist of legitimate From addresses so that the spam filter will allow through anything from these addresses. Most personal spam filters, like the one shown in Exhibit 20.3, can read your address book, and add all of the entries to your personal whitelist. New correspondents can be added over time.

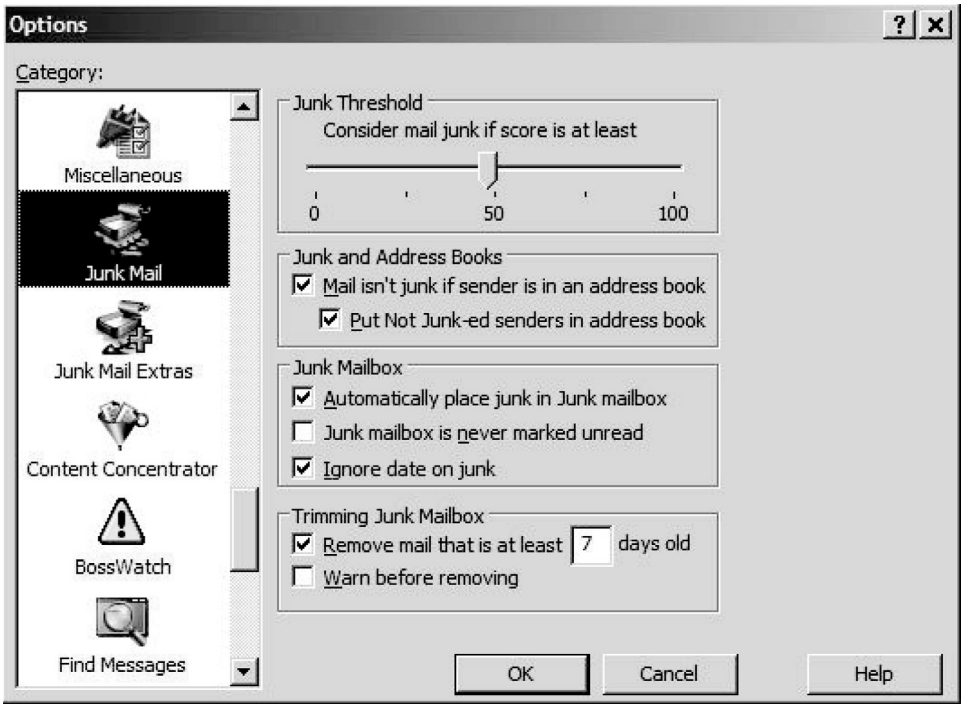


EXHIBIT 20.3 Junk Mail Filter Controls

Most personal spam filters direct the messages they identify as spam into a special folder or in basket where they can be reviewed by the user, a process sometimes referred to as quarantining. To avoid wasting hard drive space, the filtering software can be programmed to empty the quarantine folder at set intervals or simply to delete suspected spam older than a set number of days. This quarantine approach gives the user time to review and reclaim wrongly suspected spam.

20.4.5.2 ISP Filtering Faced with customer complaints about increased levels of spam around the turn of the century, ISPs began to institute spam filtering. However, they were hesitant to conduct content-based spam filtering, fearing that it might be construed as reading other people's email and lead to claims of privacy invasion. Yet ISPs must be able to read headers to route email, so filtering on the From address and Subject field was introduced (hence the increasingly inventive attempts by spammers to randomly vary Subject text). Some ISPs found that distaste for spam reached a level at which some users were prepared to accept revised terms of agreement, to allow machine reading of email content to perform more effective spam filtering. Sticky legal issues still exist for ISPs, however, especially since there is no generally accepted definition of spam and no consensus on the extent to which freedom of speech applies to email. For example, do political candidates have a right to send unsolicited email to constituents? Do ISPs have a right to block it? These are questions on which the courts of law and public opinion have yet to render a conclusive verdict.

One place where content-based spam filters are being deployed with little or no concern for legal challenges is the corporate network. Based on the fact that the company network belongs to the company, the right to control how it is used trumps

concerns over privacy. Employees do not have the right to receive, at work, whatever kind of email they want to receive. And most companies would argue that there is virtually no corporate obligation to deliver email to employees. For more details of email policies, see Chapter 48 in this *Handbook*.

20.4.5.3 Filtering Services Companies like Brightmail and Postini arose in the late 1990s to offer filtering services at the enterprise level. Brightmail developed filters that sit at the gateway to the enterprise network and filter based on continuously updated rules derived from real-time research about the latest spam outbreaks. Postini actually directs all of a company's incoming email to its servers and filters it before sending it on. With specialization, filtering services can utilize a wide range of spam fighting techniques, including:

- Block lists
- Header analysis
- Content analysis
- Filtering against a database of known spam
- Heuristic filtering for spamlike attributes
- Whitelisting through reputational schemes

Because these systems constantly deal with huge amounts of spam, they are able to refine their filters, both positive and negative, quickly and relatively effectively. And they catch a lot of spam that would otherwise reach consumers. For example, by 2008, nine of the top 12 ISPs were using Brightmail, which claims a very low false positive rate and 95 percent catch rate. Unfortunately, that still means that some spam gets through, and when spammers can direct millions of messages an hour at a network, the volume of what gets through still can pose a threat.

Another approach uses the collective intelligence of subscribers to identify spam quickly and spread the word via network connections. Cloudmark, for example, had several million subscribers at the time of writing (April 2008) paying about $40 per year to receive nearly instantaneous updates, from servers that receive and categorize reports from members on any spam that gets through. A member's credibility score rises with every correct identification of spam and sinks with incorrect labeling of legitimate email as spam (e.g., newsletters the member has forgotten subscribing to). The reliability of reporters helps the system screen false information from spammers that might try to game the system by claiming that their own junk email was legitimate.

20.4.5.4 Collateral Damage There are two major drawbacks to spam filters. First, they allow spammers to continue spamming. In other words, because they must decide, on a message-by-message basis, which messages are spam and which are not, filters consume a lot of resources, in some cases more than if all the spam was allowed through. When block lists and whitelists work well, they can substantially reduce the amount of spam that reaches the filtering stage, but eventually all filters are serial, and thus resource constrained and resource intensive.

Second, filters sometimes err, in one of two ways. They sometimes produce a false positive, flagging a legitimate message as spam, preventing a much-needed message from getting to the recipient in a timely manner. And they sometimes produce false

negatives, allowing spam messages into the inbox. False negatives and false positives impose a drag on productivity.

20.4.5.5 Snowshoe Spamming It was inevitable that spammers would adapt to the filtering measures described above. By late 2012, an increasing number of spammers were distributing their spam among a large number of compromised servers, leading to the term *snowshoe* spamming. McAfee Labs summarized the problem as follows:

> Snowshoe spamming is now one of the biggest spam problems. The issue has exploded over the past two years and will continue to increase sharply due to lack of exposure by law enforcement authorities and threats of lawsuits by companies using the illegal email lists.
>
> The phenomenon is characterized by the following:
>
> - Spammers blast out millions and millions of blatantly illegal spam messages every day from newly rented hosts until they get evicted from their subnetworks or move on.
> - Recipients have their inboxes bombarded with these spam messages and are unable to opt out of them because they are not sent from a legitimate source.
> - The result of snowshoe spamming is permanently blacklisted addresses and sometimes subnetworks.
> - Because spamming is seen as simply annoying rather than malicious, authorities have largely ignored this problem, despite the growing volumes of unwanted email originating from these sources. Companies using these shady marketers have threatened to file defamation lawsuits when researchers have tried to expose this activity.[28]

20.4.6 Network Devices The rapid rise of spam volumes in the late 1990s occurred at a time of massive email-enabled computer virus and worm attacks, arising from the classic malware motivation of bragging rights. Analysis of these problems, according to the classic factors of means, motive, and opportunity, revealed that spammers differed quite significantly from virus writers. Spammers are mostly motivated by money, not bragging rights. This insight opened up a new line of defense against spam, removing that motivation.[29]

All that remained was to understand how spammers make money, that is, the economics of spam, then to find a way to disrupt those economics. The classic spam model is to send out a vast number of emails offering a product or service, relying on the fact that at least some of these offers will reach real people, at least some of whom will make a purchase. If enough sales are made to create a profit over and above the cost of doing business, the spammer will continue to operate. Although a good public education and awareness program can reduce the number of people who buy the spammers' offers, it is unlikely to reduce that number enough. Because filtering spam out of the message stream reduces the spammer's return on investment, it tends to make spammers more inventive in their attempts to beat filters and to send even more spam in the hopes that enough will get through. In fact, spammers were able to find companies willing to sell them bandwidth on a massive scale, including dedicated point-to-point T3 digital phone line connections.

Another way to attack the economics of spam is by following the money. Every deal has to be closed, typically via a Website. Is it possible to shut down the spammer's Websites? This line of inquiry led the author to an interesting finding: Spam goes stale very quickly. An examination of spam archives showed that most links in old spam were dead, sometimes because the Website was shut down by the hosting company,

sometimes because the spammer did not want to risk being identified. Thus, the key to the economics of spam was revealed: time. If you slow spam down, spammers cannot generate enough responses before response sites are taken down.

Shortly after this realization, network security expert David Brussin devised a means of slowing down spam through TCP/IP traffic shaping. This became the technology at the heart of the antispam router. Instead of looking at each message in turn to determine if it is spam, the antispam router samples message traffic in real time and slows that traffic down if the sample suggests the traffic contains spam. To spammers, or rather the spamming software used by spammers, a network protected with an antispam router behaves as if it is on a 300-baud modem. In other words, the connection is too slow to deliver enough messages quickly enough for that one-in-a-million hit that the spam scheme needs in order to make money before the Website is shut down. The spamming software quickly drops the connection; no messages have been lost and no messages have been falsely labeled as spam. Delivery of legitimate email has not been affected. One or two spam messages have been delivered, but a lot less than are allowed through by the 95 percent catch rate of a typical spam filter.

There are two main tricks to this technology. One is tuning the TCP/IP traffic shaping, the other is tuning the sampling process. The latter needs regular updates about what spam looks like currently, so that it can identify a spammy connection as quickly as possible. These updates can be derived from services that constantly identify new spam. An application control panel is used to handle the tuning of the traffic shaping. Network administrators have found that a properly tuned antispam router can reduce bandwidth requirements by as much as 75 percent. Unfortunately, an antispam router works best on a high-volume connection, protecting MTAs at ISPs, larger companies, schools, and government agencies. There is no desktop version. Still, the technology has become a valuable part of the antispam arsenal, although the scourge of spam still continues.

20.4.7 Email Authentication On a technical level, what allows spam to continue to proliferate is the lack of sender authentication in the SMTP protocol. Several initiatives have sought to change this situation, either by changing the SMTP protocol or by adding another layer to it. One obvious way to authenticate email is via the sender's domain name, and numerous schemes to accomplish this have been proposed:

- **Sender Policy Framework (SPF)**. An extension to SMTP that allows software to identify and reject forged addresses in the SMTP MAIL FROM (Return-Path) that are typically indicative of spam. SPF is defined in Experimental RFC 4408.[30]
- **Certified Server Validation (CSV)**. A technical method of email authentication that focuses on the SMTP HELO-identity of MTAs.
- **SenderID**. An antispoofing proposal from the former MARID IETF working group that joined Sender Policy Framework and Caller ID. Sender ID is defined primarily in Experimental RFC 4406.[31]
- **DomainKeys**. A method for email authentication that provides end-to-end integrity from a signing MTA to a verifying MTA acting on behalf of the recipient. It uses a signature header verified by retrieving and validating the sender's public key through the Domain Name System (DNS).
- **DomainKeys Identified Mail (DKIM)**. A specification that merges DomainKeys and Identified Internet Mail, an email authentication scheme supported by Cisco, to create an enhanced implementation.

When properly implemented, all of these methods can be effective in stopping the kind of forgery that has enabled spam to dominate email. Each has varying pros and cons and requirements. SPF, CSV, and SenderID authenticate just a domain name. SPF and CSV can reject forgeries before message data is transferred. DomainKeys and DKIM use a digital signature to authenticate a domain name and the integrity of message content. In order for SenderID and DomainKeys to work, they must process the headers, and so the message must be transmitted.

Obviously a scheme like DomainKeys makes email more complex than SMTP. Additional data and processing effort are required. DNS records have to be expanded and maintained. However, the effort may be worth it. Although something like DomainKeys does not prevent email abuse, it does make abuse of email on protected domains easier to detect and track, and this is a deterrent to spamming now that several stiff sentences for spamming have been handed down.[32] Since 2004, the email services of both Yahoo and Google have signed outgoing email with DomainKeys, but it is not yet clear if any one of the described schemes, or some derivative thereof, will become the standard for all email. And therein lies the rub. Spam could be effectively sidelined *if* all legitimate email was reliably authenticated; recipients could simply ignore all unauthenticated messages. However, that is a large *if*.

20.4.8 Industry Initiatives Back in 2002, at the urging of consumer advocates, government agencies, and many large corporate users of email, the major email service providers started to hold meetings to discuss a unified approach to improving email and eliminating spam. There were high hopes that the problem of spam would be solved by this group, which was known in the business as AMEY, for AOL, Microsoft, Earthlink, and Yahoo!. There was certainly plenty of encouragement. Early in 2003, the FTC, America's main consumer protection agency, convened a conference on spam, and the commissioners made it clear they wanted action. Indeed, by the end of 2003, Congress had passed the first federal antispam legislation (see Section 20.4.9).

In January 2004, then Microsoft chairman and chief executive Bill Gates announced to the World Economic Forum, "Two years from now, spam will be solved." Two months later, the AMEY companies filed lawsuits against persons alleged to be major spammers, claiming that spam cost businesses in North America $10 billion each year in lost productivity, network upgrades, and destroyed or lost data. However, by January 2006, spam constituted more than three-quarters of all email. What went wrong? Repeated efforts to get the AMEY companies to adopt a new, open, royalty-free standard have run aground on concerns over intellectual property rights. Both Microsoft and Yahoo! have asserted ownership over parts of the various mechanisms put forward to authenticate email and thus solve the spam problem. Ironically, a lack of trust between the large ISPs has resulted in a lack of trust of the email they deliver.

20.4.9 Legal Remedies Laws against spam have been passed by many countries, including the United States. The CAN-SPAM Act of 2003 (Controlling the Assault of Non-Solicited Pornography And Marketing Act, effective January 1, 2004) established certain requirements for anyone sending commercial email. The law also provides penalties for spammers and for companies whose products are advertised in spam. The law addresses "email whose primary purpose is advertising or promoting a commercial product or service, including content on a Website."[33]

Enforcement of the CAN-SPAM Act is primarily the job of the FTC, but the act also gives the Department of Justice authority to enforce the act's criminal sanctions. Other federal and state agencies can enforce the law against organizations under their

jurisdiction, and ISPs can also sue CAN-SPAM violators. Here are the law's main provisions, as stated by the FTC:

- **Bans false or misleading header information**. The From, To, and routing information of your email—including the originating domain name and email address—must be accurate and identify the person who initiated the email.
- **Prohibits deceptive subject lines**. The subject line cannot mislead the recipient about the contents or subject matter of the message.
- **Requires that your email give recipients an opt-out method**. You must provide a return email address or another Internet-based response mechanism that allows a recipient to ask you not to send future email messages to that email address, and you must honor the requests. You may create a menu of choices to allow a recipient to opt out of certain types of messages, but you must include the option to end any commercial messages from the sender.

  Any opt-out mechanism you offer must be able to process opt-out requests for at least 30 days after you send your commercial email. When you receive an opt-out request, the law gives you 10 business days to stop sending email to the requestor's email address. You cannot help another entity send email to that address or have another entity send email on your behalf to that address. Finally, it is illegal for you to sell or transfer the email addresses of people who choose not to receive your email, even in the form of a mailing list, unless you transfer the addresses so another entity can comply with the law.
- **Requires that commercial email be identified as an advertisement and include the sender's valid physical postal address**. Your message must contain clear and conspicuous notice that the message is an advertisement or solicitation and that the recipient can opt out of receiving more commercial email from you. It also must include your valid physical postal address.

  Note that messages which are transactional or relationship based—that is, email that "facilitates an agreed-upon transaction or updates a customer in an existing business relationship"—are also covered in the sense that these messages may not contain false or misleading routing information.

Readers will find extensive discussion about whether CAN-SPAM has been—and could ever be—effective by searching the Web with any search engine. For a critique of the law published in February 2004, see "Can CAN-SPAM can spam?"[34]

20.5 PHISHING. Phishing is the use of unsolicited commercial email to obtain—electronically fish for—information about you, information that you would not normally disclose to a stranger, such as your bank account number, personal identification number (PIN), and other personal identifiers such as Social Security number. There was virtually no phishing activity before 2002 and relatively little in 2003, but by 2004 phishing had become an everyday threat and has continued unabated ever since. This is one category of threat that could have been squelched, along with the rest of spam, if industry leaders had chosen to cooperate rather than to compete for customers based on promising "better antispam than the other guys."

20.5.1 What Phish Look Like. Most of us first became aware of *phishing* when we received an email about a problem with an account, typically a bank account, but possibly an eBay, PayPal, Amazon, or other online account. A typical phishing

message is designed to look as if it had been sent by a large enterprise, such as the Bank of America, complete with accurate copies of the company logo, typeface, and terminology. At first glance, such messages can be quite convincing. If you receive a message like this that references an institution at which you *do* have an account, you might be tempted to read it. You might even be tempted to follow the instructions, and these are likely to lead you to a Website that asks for confidential information. Consider the text of a typical phishing message:

> You are receiving this message, due to you protection, Our Online Technical Security Service Foreign IP Spy recently detected that your online account was recently logged on from am 77.32.11.84 without am International Access Code (I.A.C) and from an unregistered computer, which was not verified by the Our Online Service Department.
>
> If you last logged in you online account on Monday May 5th 2007, by the time 6:45 pm from an Foreign Ip their is no need for you to panic, but if you did log in your account on the above Date and Time, kindly take 2-3 minute of your online banking experience to verify and register your computer now to avoid identity theft, your protection is our future medal.

What looks, at first glance, to be a very official and technical message, turns out to be full of errors. The best thing to do with these messages is delete them. The worst thing to do is respond to them. Until there is a major overhaul of email security, no reputable institution will be using email to request changes or updates to confidential account information.

If you received an email like the one discussed above but did not have a Bank of America account, you might have been confused. The fact is, the people sending these phishing messages usually have no idea whether the recipients have accounts at the institution named in the message. Indeed, this type of mismatch is the easiest way to spot some phishing messages. However, if the phisher, who has likely sent out millions of copies of the same email, gets lucky and you do happen to have an account at the named institution, or if the email is generic, then things get a little trickier. You are more likely to open a message that appears, often quite convincingly, to come from your bank. If you cannot resist looking at email about an account problem, here are some of the clues that the message is bogus. (Note that we are not implying that messages lacking these clues are therefore legitimate.)

- **Deceptive link**. Most phishing messages make an effort to look like they are legitimate, for example, by using logos and graphics stolen from the Website of the targeted institution. All phishing messages we have seen also include a link to a Website, where you are asked to provide the data the phisher is trying to steal. However, this link typically is disguised. For example, the link might be long and complicated and include the name of the bank but actually not take you to the bank's Website. Alternatively, the link may appear to be plain and simple text but in fact it is HTML-coded to go somewhere else. Some email programs, such as Eudora, will warn you of this deception and show you the real link when you place your mouse over the link text prior to clicking.

  The design of the link in a phishing scheme can be critical to its success. An increasing number of users are appropriately wary of long, complex, or merely numerical IP address links in email (e.g., "Click http://123.212.192.68"). By the use of HTML coding, messages typically obscure the destination address. However, the link is also relevant in the next stage of the attack, when the victim clicks the link. If the URL of the phishing site appears as a numerical address in the

URL field of the browser, the victim may become suspicious. Various techniques are used to make this address look plausible.

- **Change PIN/Password**. Emails that ask you to change your account access credentials are highly suspect. Legitimate companies do not make such respects via email precisely because email is so unreliable. No security update to a reputable banking Website is going to ask that you log in to your account to reset it or to prevent its being suspended. And why would a legitimate message ask you to use a password *that has supposedly been compromised*? No government agency is going to ask for your credentials or personal information via email. And no lottery on Earth uses email to notify winners. Ignore these messages, or report them as spam, and move on.
- **Bad spelling, grammar, and logic**. Whoever thought those tedious grammar lessons could be so useful? Bad grammar, spelling, and even faulty logic can be the fastest way to spot bogus email. Consider this example: "Therefore, if you are the rightful holder of the account please fill in the form below so that we can check your identy.[sic]" There is a telling typo here (*identy* for *identity*), and the logic is hopeless. Think about it: Why would a bank send an email to someone if it was not sure the person was the rightful holder of the account? Again, this is not how real companies do business today, so just move on.
- **Generic phish**. Generic account warnings are particularly nasty. They are one way that phishing attacks try to get around the problem of not knowing where the victim (you) has an account. For example, all bank accounts in America are insured by an institution called the Federal Deposit Insurance Corporation (FDIC). In 2004 someone created a particularly nasty attack that preyed on this fact, potentially snagging anyone with a bank account. This was one of the first phishing attacks to fake the linked URL using a vulnerability in Microsoft Internet Explorer to mask the actual URL. If you clicked the link in the message, the site you went to really looked as if it were www.fdic.gov but actually led to the phisher's site.

20.5.2 Growth and Extent of Phishing. Any new form of computer abuse can be said to have arrived when a toolkit facilitating the abuse becomes available. Phishing toolkits first appeared in 2006, offering scripts that enabled attackers to automatically set up phishing Websites spoofing the legitimate Websites of different brands (including the illegal appropriation of images and logos that consumers associate with those brands). These scripts helpfully generate corresponding phishing email messages. Criminals continue the evolution of their tools without respite; a January 2013 report found that spear phishers have applied a new tool called Bouncer that adapts their URLs to include unique identifiers for their intended victims; attempting to access the criminals' pages without a valid identifier in the URL results in a 404 (no such page) error, thus interfering with researchers' analysis of the phishing pages.[35]

In February 2013, analysts found that "When security experts looked into some of the highest profile hacks in recent years—one particular criminal group kept on coming to their attention. The Comment Group, which industry insiders say is based in China, offer hacking for hire—be it for individuals, corporations or governments…" They research individual companies or organizations to locate detailed information that allows highly specific topics and even content in the phishing messages. For example, a Coca-Cola executive reportedly opened a phishing email supposedly from his own boss; the link he clicked on downloaded spyware into his computer and allowed Chinese

industrial spies to extract information which stymied the acquisition of China's largest soft-drinks company.[36]

In April 2013, the U.S. Department of Homeland Security (DHS) warned "organizations that post a lot of business and personal information on public web pages and social media sites" not to do so. In October 2012, phishers harvested detailed information about employees from a public posting by an energy company listing attendees at a conference. In addition to spear-phishing attacks on named individuals, "Malicious emails that appeared to be from one of the attendees were sent to others on the list informing them of a change in the sender's email address. Recipients were politely asked to click on an attached link that promptly took them to a site containing malware."[37]

The APWG is a valuable repository of statistical information about phishing.[38] Their report on phishing during the fourth quarter of 2013 includes the following findings (these are direct quotations formatted as bullet points with page references removed):

- Phishing attacks against online game players saw a massive increase, from 2.7 percent of all phishing attacks in Q3 to 14.7 percent in Q4.
- Financial services continued to be the most targeted industry sector in the fourth quarter of 2012, with payment services close behind.
- Online gaming credentials are valuable to certain criminals, who sell them on the black market. In-game items held in those accounts can also be sold by phishers for real-world cash. Victims can even have their real-life identities stolen.
- Attacks against social media doubled to 6 percent, up from 3 percent in the third quarter.
- During Q4, about 30 percent of personal computers worldwide were infected with malware.
- More than 57 percent of PCs in China may have been infected, while PCs in European nations were infected least often.
- Except for October 2012, the number of phishing sites declined every month from April 2012 through December 2012.
- April 2012 saw 63,253 unique phishing sites detected, falling to 45,628 in December 2012.
- The APWG received reports of 28,195 unique phishing sites in December. December's total was 31 percent lower than the high of 40,621 reports in August 2009.
- Use of crimeware dipped slightly in this quarter from the previous, as did the use of data-stealing malware.
- The use of other malware has increased by a statistically significant amount from the previous quarter.[39]

20.5.3 Where Is the Threat? Phishing appears at first to be a personal problem, a matter of consumer computer security rather than enterprise information security. Indeed, the federal agency charged with consumer protection, the FTC, has been active in educating consumers about the problem and the ways to avoid getting scammed by these messages.[40] However, phishing is also a computer-based threat to the information and welfare of companies as well as a threat to e-commerce in general.

Phishing attacks collect usernames and passwords. Many people use the same credentials for work-related systems as they do for personal systems (including the head of a highly secret government facility who was found to be using her bank ATM PIN as her top-secret network password). Criminal hackers are known to have penetrated systems by harvesting personal credentials and applying them to the target's business login. So a company computer security awareness program would do well to include warnings about phishing attacks.

Beyond the penetration threat, phishing can undermine consumer trust in a company. Although it hardly seems fair for consumers to resent Bank of America because a criminal has attempted to defraud them by abusing the bank's identity, such resentments do exist, and they need to be understood in a business context. Automation improves the profitability of banking, which is why banks encourage customers to use online services and offer incentives to drop paper statements and notifications. However, if banks are perceived as providing automation on the cheap, with insufficient safeguards built in to protect customer privacy and account access, some consumers will object, slowing down the pace of automation and jeopardizing productivity increases, potentially impacting the bottom line.

20.5.4 Fighting Phishing Some clever technological defenses specific to phishing have been put forward, such as methods for enabling browsers to verify URLs, but the best defense is a twofold approach that is entirely obvious. The first step is education, at the consumer and corporate level, teaching people how to spot phishing attacks and avoid falling victim to them. The second step is fixing the fundamentals of email in the ways outlined earlier with respect to spam. From a technical perspective, we already have the technology to do this. All that is missing is a willingness on the part of the leading service providers to take concerted action. Perhaps they could be encouraged to do so by the banks and other institutions that stand to gain a great deal if email is made more secure.

Failure to act on email security has already cost billions of dollars in wasted resources and lost productivity, but the effects go further than that, as evidenced by the emerging relationship among phishing, spamming, criminal hacking, Trojan code, personal data harvesting, and account compromise. The application of skills such as virus and worm writing to commercial ends has been enabled by the willingness of spammers to pay for compromised hosts with which to launch attacks and collect dollars and/or data. Spamming techniques have enabled the growth of phishing, which in turn has solidified the black market in compromised hosts and purloined personal data. New types of attack are constantly emerging from this unholy alliance between coders and criminals. Email addresses harvested by spam may be used for spear phishing attacks, where a specific company is targeted through emails sent to known customers or employees (the Department of Defense faced a rash of spear-phishing attacks in 2006). Trojan code distributed by spam methods has been used to corrupt local domain name servers and produce the same effect as a phishing message: Users are misdirected to malicious Websites simply by clicking on an apparently valid URL. (DNS attacks of this type are referred to as pharming.)

20.6 TROJAN CODE. Like the original Trojan horse, deployed by the Greeks to defeat the Trojans who were protected by the impregnable walls of the city of Troy, a computer Trojan is a bad thing disguised as a good thing (where Troy is your computer and the Greeks are the people who would like to gain unauthorized access). The technology has come a long way from the "steed of monstrous height" described



EXHIBIT 20.4 Screensaver Trojan

in Virgil's *Aeneid,* but the goal of Trojan code is the same as that of the original Trojan horse: Fool the defender of a protected place into granting access to outsiders. (See Section 15.3.2.17 for more on the origins of the name.)

Almost as soon as the Internet provided a way to distribute executable code, either as download links or email attachments, some people decided to exploit this ability to distribute their code without explicit permission. In other words, an innocent party might be tempted to download an apparently innocuous executable, which was actually something else. The person doing the downloading does so intentionally but is ignorant of the intentions of the person who crafted the Trojan code within the executable.

20.6.1 Classic and Recent Trojans. Despite the fact that the screens of today's computers do not need saving, screensavers continue to be a source of Trojan code. Some users apparently find the promise of animated waterfalls and fish tanks hard to resist. This leads users to download screensavers that they encounter on Websites or open screensaver files that they receive in email. In Exhibit 20.4 you can see an example of one screensaver.

The creator of this Trojan screensaver, discovered in February 2007, was apparently unhappy with a file-sharing network popular in Japan and named Winny. When this screensaver is opened, the Trojan displays an image warning the computer operator not to use Winny. This tactic is reminiscent of some of the very earliest viruses, which sought to spread relatively innocuous messages rather than cause damage. In fact, code does not have to have a malicious intent to be considered a Trojan—it merely needs the intent to get installed without permission. Unfortunately, in this particular case, the Trojan does go on to destroy data by overwriting certain files (e.g., those with .txt and .jpg extensions).

Another example of a screensaver Trojan is the one reported to be circulating as an email attachment called bsaver.zip in July 2007. When the file within the ZIP attachment is opened, the system is infected with the Agent-FZB Trojan horse, which then drops two rootkits to evade detection by security software and to make the system accessible to unauthorized users. (See Chapter 15 in this *Handbook* for more on rootkits.) This

screensaver was distributed by a spam campaign with message subject lines such as *Life is beautiful*, *Life will be better*, *Good summer*, and *help you*. The message text included phrases like "Good morning/evening, man! Really [sic] cool screensaver in your attachment!"

More recent reports about Trojans include the following cases and studies:

- The BackDoor.Wirenet.1 Trojan was identified in August 2012; the malware "is the first Trojan Horse program that works on the Mac OS X and Linux platforms that is 'designed to steal passwords stored by a number of popular Internet applications.' "[41]
- The "PandaLabs Q1 Report" for 2013 found that "Trojans set a new record, causing nearly 80 percent of all computer infections worldwide. Despite their inability to replicate, Trojans are capable of triggering massive infections through compromised Web sites that exploit vulnerabilities in browser plug-ins like Java, Adobe Reader, etc. This attack method allows hackers to infect thousands of computers in just a few minutes with the same Trojan or different ones, as attackers have the ability to change the Trojan they use based on multiple parameters such as the victim's location, the operating system used, etc."[42]
- In March 2013, Kaspersky Labs reported that a new spyware attack on Tibetan freedom activists used a Trojan designed for the Android mobile-phone operating system.[43]
- The Flashback Trojan had infected more than 600,000 Macintosh computers by early April 2013 by exploiting a flaw in Java.[44]
- In April 2013, criminals sent out invitations to watch video footage of the Boston Marathon bombings. A remote-access Trojan was installed as a "WinPcap Packet Driver (NPF)" to evade notice.[45]
- In May 2013, Graham Cluley of Sophos reported on a Trojan (Mal/BredoZp-B) circulated in an email supposedly from Tiffany & Co. that claimed to include details of an export license and a payment invoice.[46]

A Trojan tries to look useful or interesting so that users will install it, but buried within is unauthorized, undocumented code for unauthorized functions, just some of which are listed here:

- Deleting files or folders or entire drives
- Changing data, in subtle or dramatic ways
- Encrypting data (possibly for purposes of extortion)
- Copying data to other computers (possibly for purposes of industrial espionage)
- Downloading files without user consent (possibly for purposes of illicit e-commerce, illegal file sharing, pornography site hosting, cheap storage)
- Corrupting browser software and network files to redirect user from legitimate Websites to fake, bogus, malicious sites
- Enabling remote access to the compromised system (sometimes referred to as a RAT, for Remote Access Trojan), often used to turn computers into zombies that can be aggregated into botnets (for the purposes of spamming, phishing, and executing denial-of-service attacks)
- Aiding the spread of viruses using dropper code to cause infections

- Disabling antivirus and firewall programs
- Disabling competing forms of malware
- Logging keystrokes to obtain data such as passwords and credit card numbers
- Reporting browsing activity (enabling spyware)
- Harvesting email addresses for the purpose of spamming and other activities such as phishing

The motives behind Trojan code typically fall into one or more of three categories: malice, bragging rights, and financial gain. The gains can be realized in several ways. Purloined data can be sold. Access to compromised machines can be sold. Encrypted or stolen data can be ransomed. Denial-of-service threats can be used for extortion. Some companies have even used Trojans to try to increase the sales and installed base of their software.

20.6.2 Basic Anti-Trojan Tactics. Unfortunately, Trojans can be tough to combat. The first line of defense is well-educated users who know better than to execute code of dubious origin. Warnings to this effect should be part of any computer security awareness program, and a computer security awareness program should be part of every enterprise security model. Technical measures can be used to combat Trojans, but because none of these is perfect, it makes no sense to rely on them to the neglect of end user education. In fact, thinking of end users as *computer operators* rather than *computer users* might be a good place to start, given that the role of user implies a relatively passive role in maintaining the integrity and health of a computer, whereas operator reflects more accurately the level of responsibility required of anyone employing a computer in either work or recreation. (A computer user might be likened to someone who merely drives a car and never checks the oil or the tires; an operator is closer to a commercial driver who knows that getting safely and reliably from A to B takes a lot more than just holding the steering wheel and pressing the correct pedal.)

Technical measures against Trojans start with keeping all operating system and application patches up to date, universal use of memory-resident antivirus software that scans incoming files, and regular scanning of the entire system against a regularly updated virus database. At the same time, it helps to run memory-resident antibot software, such as Norton AntiBot, designed to recognize and thwart activity indicative of botnets. This software does not directly prevent Trojans from getting onto a computer, but it minimizes the damage that a Trojan can cause.

A good antispam solution is also necessary because spam is a major vector for Trojan attacks. If Jim in accounting never gets that message about the cool screensaver, he will not ever be tempted to open the attached file. It is not just screensavers that can be tempting. In April 2007, spam techniques were used to distribute, on a massive scale, messages with subject headings like "Worm Alert!" and "Worm Detected." The messages came with a ZIP file attachment that posed as a patch that would prevent the bogus attack. When recipients, alarmed into action, went to open the ZIP file, they found that it was password protected, and some took that as a sign it was genuine. (The password was included in the message.) Proceeding with installation led to a rootkit, disabling of security software, theft of confidential information from the affected machine, and enrollment in a botnet of compromised computers. In one 24-hour period, Postini counted nearly 5 million copies of this Trojan spam directed at users of its antispam service. The company calculated that this one spam accounted for 87 percent of all malware being spread through email during that time.

20.6.3 Lockdown and Quarantine. More drastic anti-Trojan measures include preventing unauthorized changes or additions to executables and preventing systems from connecting to networks until they have been vetted. The idea of defeating malicious executables by freezing the operating system, and authorized applications in a known good state, goes back a long way, at least to the early days of antivirus efforts. However, the complexity of most operating systems and application code today makes such an approach challenging at best. Consider the emergence of Patch Tuesday as a standard means of maintaining the world's most widely used operating system. Indeed, for many years now, a lot of software development has been predicated on the assumption that patches and updates can always be pushed out if needed, arguably leading to far less rigor in coding practices than when production code was burned to disk, at considerable expense, and changes required further shipments of disks.

The idea of quarantining computers when they attempt to connect to a network, using some form of network access control, offers a slightly different approach to the idea of locking down machines so that unauthorized code cannot be installed. Some enterprise network security managers have pursued this approach because it is increasingly difficult to control some network end points, such as the company laptop that travels with the employee to client sites and conferences, hotels, and WiFi hot spots, even spending time in the employee's car and home. Laptops used in this manner face any number of attack vectors. Traditional defense mechanisms can be applied to these end points, including password protection, biometric authentication, disk encryption, and antivirus programs. However, these may not be enough, as many enterprises have discovered to their cost. So why not prevent these machines from connecting to the enterprise network until they have been scanned for unauthorized executables, like a routine health check? This approach holds promise but is not easy to implement. Furthermore, it leaves machines open to attack while they are away from the corporate network. A spam-distributed phishing exercise or screensaver Trojan could still lobotomize an end point between checkups and compromise confidential personal or corporate data. As you can see from Exhibit 20.5, a table compiled by Symantec in 2011 to show prices paid for various forms of data in the underground economy in 2009 and 2010, the financial incentives for such activities are real.

EXHIBIT 20.5 Table of Prices Paid for Data Traded in the Underground Economy[47]

Overall Rank			Percentage		
2010	**2009**	**Item**	**2010**	**2009**	**2010 Price Ranges**
1	1	Credit card information	22%	19%	$0.07–$100
2	2	Bank account credentials	16%	19%	$10–$900
3	3	Email accounts	10%	7%	$1–$18
4	13	Attack tools	7%	2%	$5–$650
5	4	Email addresses	5%	7%	$1/MB–$20/MB
6	7	Credit card dumps	5%	5%	$0.50–$120
7	6	Full identities	5%	5%	$0.50–$20
8	14	Scam hosting	4%	2%	$10–$150
9	5	Shell scripts	4%	6%	$2–$7
10	9	Cash-out services	3%	4%	$200–$500 or 50%–70% of total value

20.7 CONCLUDING REMARKS. The threats addressed in this chapter are still evolving, as are the countermeasures discussed here, although some of those mentioned may be rendered obsolete by new developments. If there is one lesson security professionals can learn from studying the threats discussed herein, it is that continued failure to improve the underlying security of Internet email will prolong the onslaught of spam, phishing attacks, and Trojans. The second lesson is that we must continue to educate computer users in order to prevent them falling prey to the deceptions of computer abusers. A third lesson might well be that more effective prosecution of computer criminals is needed, along with stiffer sentences.

20.8 FURTHER READING

Brown, B. C. *The Complete Guide to Email Marketing: How to Create Successful, Spam-Free Campaigns to Reach Your Target Audience and Increase Sales*. Ocala, FL: Atlantic Publishing Group, 2007.

Goodman, D. *Spam Wars: Our Last Best Chance to Defeat Spammers, Scammers & Hackers*. New York, NY: Select Books, 2004.

Haskins, R., and D. Nielsen. *Slamming Spam: A Guide for System Administrators*. New York, NY: Addison Wesley Professional, 2004.

Jackobsson, M., and S. Myers, eds. *Phishing and Countermeasures: Understanding the Increasing Problem of Electronic Identity Theft*. Hoboken, NJ: John Wiley & Sons, 2006.

James, L. *Phishing Exposed*. Rockland, MA: Syngress, 2005.

Lininger, R., and R. D. Vines. *Phishing: Cutting the Identity Theft Line*. Indianapolis, IN: John Wiley & Sons, 2005.

McWilliams, B. S. *Spam Kings: The Real Story behind the High-Rolling Hucksters Pushing Porn, Pills, and %* @)# Enlargements*. Sebastopol, CA: O'Reilly, 2004.

Schryen, G. *Anti-Spam Measures: Analysis and Design*. Berlin: Springer, 2007.

Silver Lake, eds. *Scams & Swindles: Phishing, Spoofing, Spyware, Nigerian Prisoner, ID Theft*. Aberdeen, WA: Silver Lake Publishing, 2006.

Spammer-X. *Inside the SPAM Cartel: Trade Secrets from the Dark Side*. Rockland, MA: Syngress, 2004.

Wolfe, P., C. Scott, and M. Erwin. *Anti-Spam Tool Kit*. Emeryville, CA: McGraw-Hill Osborne Media, 2004.

Zdziarski, J. *Ending Spam: Bayesian Content Filtering and the Art of Statistical Language Classification*. San Francisco, CA: No Starch Press, 2005.

20.9 NOTES

1. U.S. Internal Revenue Service, "Report Phishing," IRS Website, May 8, 2013, www.irs.gov/uac/Report-Phishing
2. Symantec, "2013 Internet Security Threat Report, Volume 18," *Symantec | Security Response Publications*, April 15, 2013, www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v18_2012_21291018.en-us.pdf
3. Dean Turner, "Symantec Internet Security Threat Report, Volume X," *Symantec | White Papers*, September 22, 2006, http://eval.symantec.com/mktginfo/enterprise/white_papers/ent-whitepaper_symantec_internet_security_threat_report_x_09_2006.en-us.pdf.[47] The diagram in Exhibit 20.1 and accompanying text originally appeared in *Trusted Email Open Standard: A Comprehensive Policy and Technology Proposal for Email Reform*, a white paper published in 2003 by a company called

ePrivacy Group; contributors to the paper included Vincent Schiavone, David Brussin, James Koenig, and Ray Everett-Church.

4. The diagram in Exhibit 20.1 and accompanying text originally appeared in *Trusted Email Open Standard: A Comprehensive Policy and Technology Proposal for Email Reform*, a white paper published in 2003 by a company called ePrivacy Group; contributors to the paper included Vincent Schiavone, David Brussin, James Koenig, and Ray Everett-Church.
5. Karl Taro Greenfeld, "How Spam Meat Has Survived Spam E-Mail," *Bloomberg Businessweek | Lifestyle*, May 17, 2012, www.businessweek.com/articles/2012-05-17/how-spam-meat-has-survived-spam-e-mail
6. Monty Python, "Spam," *Youtube | Monty Python Channel*, uploaded January 13, 2009, www.youtube.com/watch?v=M_eYSuPKP3Y&list=UUGm3CO6LPcN-Y7HIuyE0Rew&index=32
7. See online article at www.templetons.com/brad/spamterm.html
8. Austin Modine, "Virginia De-convicts AOL Junk Mailer Jeremy Jaynes: Overturns Anti-Spam Law, Invokes Founding Fathers." *The Register*. September 13, 2008, www.theregister.co.uk/2008/09/13/virginia_overturns_antispam_conviction
9. Grant Gross, "Court Overturns Virginia Spam Law, Conviction," *PCWorld*, September 12, 2008, www.pcworld.com/article/151014/spam_law_overturned.html
10. Associated Press, "Sanford Wallace: Facebook Wins $711 Million in Case against 'Spam King,' " *Huffington Post | Tech*, October 30, 2009, www.huffingtonpost.com/2009/10/30/sanford-wallace-facebook-_n_339703.html
11. FBI, "Sanford Wallace Indicted for Spamming Facebook Users: Self-proclaimed 'Spam King' Sent More Than 27 Million Spam Messages," FBI | San Francisco Division | US Attorney's Office | Northern District of California, August 04, 2011, www.fbi.gov/sanfrancisco/press-releases/2011/sanford-wallace-indicted-for-spamming-facebook-users
12. San Jose Federal Court, "Calendar for Judge Edward J. Davila," U.S. Courts, May 23, 2013, www.cand.uscourts.gov/CEO/cfd.aspx?7143
13. The term "Email Service Provider" applies to companies that provide email addresses but not necessarily the Internet connection used to access those addresses. Google is a prime example with its Gmail; other examples include AOL, MSN, and Yahoo!, although all three of these have at times offered Internet connectivity as well as email and thus operated as ISPs as well as Email Service Providers.
14. Spam constituted approximately 90 percent of all email sent in 2007, according to numbers compiled by a variety of email service providers, including Postini, which is owned by Google.
15. See www.messagelabs.com/intelligence.aspx for the list of monthly and annual reports; the 2007 report was available from www.messagelabs.com/mlireport/MLI_2007_Annual_Security_Report.pdf
16. Robert McMillan, "90 percent of e-mail is spam, Symantec says," *Computerworld | Applications*, May 26, 2009, www.computerworld.com/s/article/9133526/90_percent_of_e_mail_is_spam_Symantec_says
17. Amar Toor, "Symantec Says 92-Percent of All E-mail Is Spam, Phishing Attacks Declining," *SWITCHED*, August 13, 2010, www.switched.com/2010/08/13/symantec-says-92-percent-of-all-e-mail-is-spam-phishing-attacks

18. Nicole Henderson, "Symantec Report Finds Spam Accounts for 73 Percent of June Email," *Web Host Industry Review*, June 28, 2011, www.thewhir.com/web-hosting-news/symantec-report-finds-spam-accounts-for-73-percent-of-june-email
19. Lance Whitney, "Spam Sinks to Lowest Level in Almost Three Years, Says Symantec: The Amount of Spam around the Globe Now Accounts for 70 Percent of All E-mail, a Sharp Decline from 2009 When It Accounted for 90 Percent," *c|net | News | Security & Privacy*, December 7, 2011, http://news.cnet.com/8301-1009_3-57338317-83/spam-sinks-to-lowest-level-in-almost-three-years-says-symantec
20. Kaspersky Lab, "Spam in 2012: Continued Decline Sees Spam Levels Hit 5-year Low," Kaspersky Lab, January 13, 2013, www.kaspersky.com/about/news/spam/2013/Spam_in_2012_Continued_Decline_Sees_Spam_Levels_Hit_5_year_Low
21. Darya Gudkova and Tatyana Shcherbakova, "Spam in January 2013," *Securelist | Analysis*, February 21, 2013, www.securelist.com/en/analysis/204792282/Spam_in_January_2013
22. Sara Radicati and Quoc Hoang, "Email Statistics Report, 2012–2016: Executive Summary," Radicati Group, April 23, 2012, www.radicati.com/wp/wp-content/uploads/2012/08/Email-Statistics-Report-2012-2016-Executive-Summary.pdf
23. Justin M. Rao and David H. Reiley, "The Economics of Spam," *Journal of Economic Perspectives* 26, no. 3 (2012): 87–100, www.aeaweb.org/articles.php?hs=1&fnd=s&doi=10.1257/jep.26.3.87
24. Located at www.networkworld.com/spam/index.jsp
25. Anti-Phishing Working Group, "About APWG," APWG Website, 2013, http://apwg.org/about-APWG
26. Janusz Lukasiak, "Blocking Spam Relaying and Junk Mail," Jisc Website, October 1999, www.jisc.ac.uk/publications/publications/blockingspamfinalreport.aspx
27. The exact terminology for block lists and black holes is a subject of ongoing debate involving possible issues of trademark infringement and political correctness.
28. McAfee, "Snowshoe Spamming Emerges as Threat to Email Security," *McAfee | Business Home | Security Awareness*, December 27, 2012, www.mcafee.com/us/security-awareness/articles/snowshoe-spamming-biggest-problem.aspx
29. David Brussin, Stephen Cobb, Ray Everett-Church, and Vincent J. Schiavone, "Network Resource Theft Prevention: Destroying the Economics of Spam," ePrivacy Group, November 2003, http://cobbassociates.com/library/spamsquelcher_wp2.pdf
30. See www.rfc-archive.org/getrfc.php?rfc=4408
31. See www.rfc-archive.org/getrfc.php?rfc=4406
32. Sentences for spam convictions have been rising steadily. Although many of those convicted have appealed, at least half a dozen have now served jail time.
33. Federal Trade Commission, "The CAN-SPAM Act: Requirements for Commercial Emailers," www.business.ftc.gov/documents/bus61-can-spam-act-compliance-guide-business
34. M. E. Kabay, "Can CAN-SPAM can spam?" *Networkworld*, February 02, 2004, www.networkworld.com/newsletters/sec/2004/0202sec1.html
35. Paul, "New Phishing Toolkit Uses Whitelisting To Keep Scams Alive," *VERACODE | the security ledger*, January 16, 2013, https://securityledger.com/new-phishing-toolkit-uses-whitelisting-to-keep-scams-alive

36. Dave Lee, "The Comment Group: The Hackers Hunting for Clues about You," *BBC | News | Business*, February 11, 2013, www.bbc.co.uk/news/business-21371608
37. Jaikumar Vijayan, "DHS Warns of Spear-Phishing Campaign against Energy Companies: Attackers Used Information from Company Website to Craft Attacks," *Computerworld | Security | Malware and Vulnerabilities*, April 5, 2013, www.computerworld.com/s/article/9238190/DHS_warns_of_spear_phishing_campaign_against_energy_companies
38. Anti-Phishing Working Group, "APWG Phishing Attack Trends Reports," Anti-Phishing Working Group Website, April 24, 2013, http://apwg.org/resources/apwg-reports
39. Greg Aaron, "Phishing Activity Trends Report: 4th Quarter 2012," Anti-Phishing Working Group Website, April 24, 2013, http://docs.apwg.org/reports/apwg_trends_report_Q4_2012.pdf
40. See www.ftc.gov and http://onguardonline.gov/phishing.html
41. Anthony Wing Kosner, "New Trojan Backdoor Malware Targets Mac OS X and Linux, Steals Passwords and Keystrokes," *Forbes | Tech*, August 31, 2012. www.forbes.com/sites/anthonykosner/2012/08/31/new-trojan-backdoor-malware-targets-mac-os-x-and-linux-steals-passwords-and-keystrokes
42. PandaLabs, "PandaLabs Q1 Report: Trojans Account for 80% of Malware Infections, Set New Record," *Panda Security | Press Room | News*, May 3, 2013, http://press.pandasecurity.com/news/pandalabs-q1-report-trojans-account-for-80-of-malware-infections-set-new-record
43. Sean Gallagher, "First Targeted Attack to Use Android Malware Discovered: Kaspersky Uncovers Trojan Spread by "Spear-phish" to Tibet Activists," *ars technica*, May 26, 2013, http://arstechnica.com/security/2013/03/first-targeted-attack-to-use-android-malware-discovered
44. Dwight Silverman, "More than Half a Million Macs Infected with Flashback Trojan Malware," *Chron | TechBlog*, April 5, 2012, http://blog.chron.com/techblog/2012/04/more-than-half-a-million-macs-infected-with-flashback-trojan-malware
45. Graham Cluley, "Sick Malware Authors Exploit Boston Marathon Bombing with Trojan Attack," *Sophos | nakedsecurity*, April 17, 2013, http://nakedsecurity.sophos.com/2013/04/17/malware-boston-marathon-bombing
46. Graham Cluley, "Breakfast Malware at Tiffany's? Trojan Horses Spammed out Widely," *Sophos | nakedsecurity*, May 22, 2013, http://nakedsecurity.sophos.com/2013/05/22/tiffany-malware
47. Symantec, "Fraud Activity Trends," *Symantec | Enterprise | Security Response | Internet Security Threat Report*, 2010, www.symantec.com/threatreport/topic.jsp?id=fraud_activity_trends&aid=underground_economy_servers

CHAPTER 21

WEB-BASED VULNERABILITIES

Anup K. Ghosh, Kurt Baumgarten, Jennifer Hadley, and Steven Lovaas

21.1 INTRODUCTION 21·1

21.2 BREAKING E-COMMERCE SYSTEMS 21·1

21.3 CASE STUDY OF BREAKING AN E-BUSINESS 21·2

21.4 WEB APPLICATION SYSTEM SECURITY 21·5

21.5 PROTECTING WEB APPLICATIONS 21·6

21.6 COMPONENTS AND VULNERABILITIES IN E-COMMERCE SYSTEMS 21·8

21.6.1 Client-Side Risks 21·8

21.6.2 Network Protocol Risks 21·10

21.6.3 Business Application Logic 21·12

21.6.4 CGI Script Vulnerabilities 21·14

21.6.5 Application Subversion 21·15

21.6.6 Web Server Exploits 21·16

21.6.7 Database Security 21·19

21.6.8 Platform Security 21·21

21.7 SUMMARY 21·22

21.8 FURTHER READING 21·23

21.9 NOTES 21·23

21.1 INTRODUCTION. This chapter systematically reviews the primary software components that make up Web applications, with a primary focus on e-commerce, and provides an overview of the risks to each of these components.[1] The goal of this chapter is to point out that every system will have risks to its security and privacy that need to be systematically analyzed and ultimately addressed.

21.2 BREAKING E-COMMERCE SYSTEMS. To make a system more secure, it may be advisable to break it. Finding the vulnerabilities in a system is necessary in order to strengthen it, but breaking an e-commerce system requires a different mind-set from that of the programmers who developed it. Instead of thinking about developing within a specification, a criminal or hacker looks outside the specifications.

Hackers believe that rules exist only to be broken, and they always use a system in unexpected ways. In doing so, they usually follow the path of least resistance. Those areas perceived to provide the strongest security, or the most resistance to hacking, will likely be ignored. For example, if a system uses Secure Sockets Layer (SSL) to encrypt Web sessions between Web clients and the Web server, a hacker will not try to

break the encryption stream but instead will look for an easier way to get at the data after they are decrypted and stored in the clear.

Hackers go where the money is—sometimes literally, sometimes not. They typically try to hack into a site only if there is some reward for their effort. Sometimes hackers are motivated by money, but as often the motivation is the lure of fame, notoriety, or acceptance by a peer group. The level of protection should be commensurate with the value of the resource being protected. For instance, a Website that publishes the menus for local restaurants may not be seen as a target for a denial of service or any other type attack. Such a Website simply is not as attractive a target as a bank's online Website, where a hacker can certainly gain in notoriety and even profit financially. Similarly, most people do not bother encrypting email messages due to lack of security awareness, and because most potential snoopers are not interested in ordinary personal email. Sensitive email from a high-profile organization should be encrypted, however, in order to protect valuable intellectual capital.

It is important to remember that the e-commerce system is a chain that is only as strong as its weakest link.[2] Hackers naturally attempt to attack that point of least resistance. This explains why a site may deliberately set up a honeypot (a sacrificial system) with appealingly vulnerable services in order to track and monitor potential hackers. In e-commerce, or any systems with a Web presence, cryptography often is perceived to provide the strongest level of protection; thus, hackers generally attack host-side services and client-side content.

Maintaining strong host security, both inside and outside the perimeters, is critically important. One unfortunate side effect of corporate perimeters is that system administrators tend to overlook host-based security, which may leave the host completely vulnerable. The result is that once hackers make it through or around perimeter devices and policies (e.g., firewalls, routers, or even receptionists), they can leverage the internal trust relationships to compromise many resources, including workstations, servers, and phone systems. The prudent administrator will exercise equal concern both at the entry to systems and within those systems.

21.3 CASE STUDY OF BREAKING AN E-BUSINESS. Consider an online investing e-business application and how a hacker might go about disassembling its security for malicious or financial gain. Online investing is very popular for several reasons. Rather than waiting for quarterly statements in the mail or dealing with a phone menu, customers can quickly view the status, balances, and current value of investment holdings by visiting the Web pages of their portfolio managers. If they wish to buy and sell equity shares on demand, they can establish online Web-enabled brokerage accounts. Exhibit 21.1 shows a simplified workflow diagram of an online investing application that enables users with established accounts to view portfolio holdings and cash balances, to update the stocks tracked, and to conduct online trades.

To see how this application can be broken, it is helpful to look at a sample network architecture that implements the online application. Exhibit 21.2 shows the network architecture of the system that implements the online investing application, along with example exploits. The system consists of the end users' client machines, the Internet, routers, firewall, front-end Web and email servers, application servers, databases, and workstations.

There are many ways a hacker could break this online application. Exhibit 21.2 shows one possible scenario. In step 1 of the attack, the hacker uses the Internet and a Web browser to misuse one of the Common Gateway Interface (CGI) scripts that invoke the application on a server. The CGI script could be a development CGI script

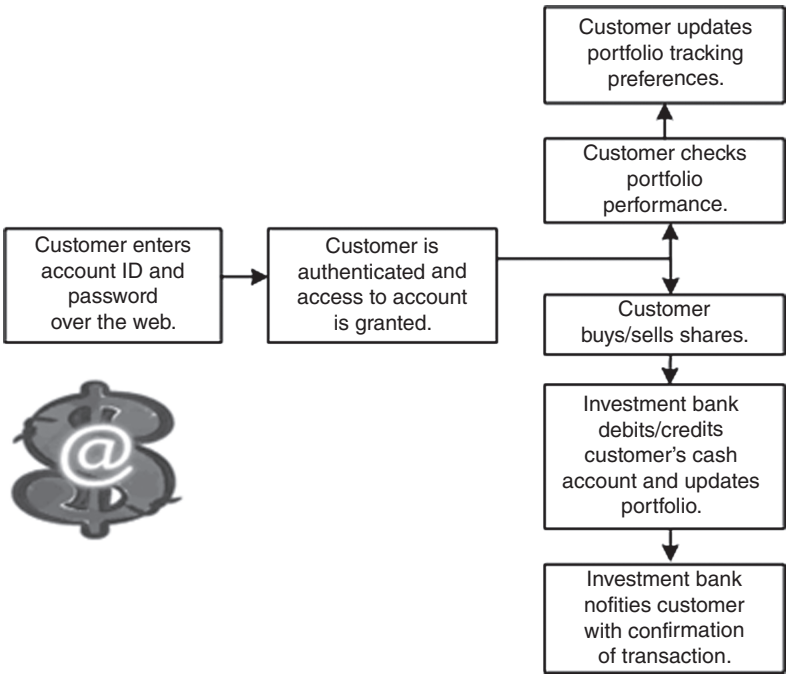


EXHIBIT 21.1 Online Investing Application

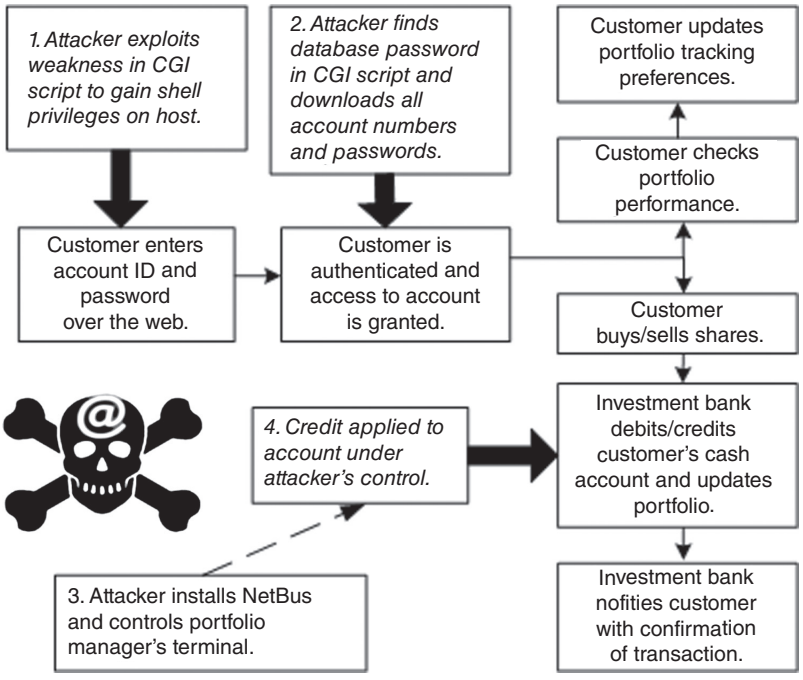


EXHIBIT 21.2 Breaking an E-Business

inadvertently left on the server before going into production, a default CGI script from an application server distribution, or a script that implements flawed logic in the online investment application. Exploiting CGI scripts via cross-site scripting or other common exploit methods allow hackers to gain shell access to Web servers. CGI script vulnerabilities are discussed later in this chapter.

The vulnerability need not be in a CGI script. Application servers can be implemented in Java, C, C++, Perl, or Python in various application server frameworks. The difficulty lies not in which language the business application logic is developed; more important are the vulnerabilities introduced by the complex logic at this middleware layer. One of the key problems in the development of application middleware is poor input sanity checking; that is, the developers fail to impose limits on acceptable input. The hacker can exploit the lack of input sanity checking to feed the application server unexpected input used in system commands. This technique can gain shell privileges on the machine.

Although application server misuse is a common way of breaking into systems, there are many other ways to gain the initial access in step 1 of the attack. For instance, the Web and mail servers may be running any of several network services, such as *FTP* and *BIND/DNS,* which may be misconfigured and thus "poisoned." The Web and mail server software themselves may be vulnerable to attack. Many popular commercial Web and mail servers have been vulnerable to buffer overflow attacks that often permit full system root privileges on the host.[3] Once attackers gain system privileges on an internal host, they can exploit the inherent internal trust often woven among machines through network policies in order to gain access to other systems on the network. This strategy is precisely what the attacker follows in step 2 of the attack illustrated in Exhibit 21.2

Once attackers have access to the various file systems on the application server, they can view source code of CGI scripts or other application middleware to discover customer account numbers, passwords, and even database administrator passwords for accessing the back-end databases. From there, they can download important and confidential client information stored in the database. In step 3 of the attack, the attacker leverages the internal privileges gained to plant backdoors into the system unnoticed. A suite of software, commonly known as a rootkit and available to hackers, allows them not only to get into a system unnoticed but also to erase their tracks in audit logs. In the example shown, the hacker installs a rogue remote administration program known as Back Orifice, which provides the ability to remotely administer the host and network with the same privileges and power as an authentic system administrator.

At this point in the attack, the hacker has assumed total control of the systems and any Web-facing or commerce-related processes, with many options including:

- Stealing customer or company information for the purpose of financial gain
- Defacing the Web pages for notoriety or to publicize an agenda
- Working in a stealthy manner to uncover proprietary business information and other confidential intellectual capital (espionage)
- Blackmailing the business with threats of discrediting it
- Subverting the application for any other personal gain

Step 4 of the attack illustrates the last case, where the attacker credits a personal cash account. The hacker must move quickly enough to withdraw these funds before traditional back-end auditing discovers the discrepancy. There have been many defaced

Web pages of government agencies, and other important sites, and many reported instances of the other cases. Of course, some companies are understandably reluctant to publicize events that might lessen customer confidence, yet in many cases, legislation now requires the disclosure of a breach.

Unfortunately, it takes only a single flawed, unpatched computer, or overlooked vulnerability, for a hacker to compromise a system as a whole. Although *defense-in-depth* (using multiple forms of security, such as firewalls on the perimeter and intrusion detection inside the network) is a popular strategy, often multiple layers of defense fall like a house of cards when a single hole is exploited. For example, an attacker who gains *root* capability can disable all other security measures. The problem is known as an *asymmetric* attack because it is much more difficult and costly to defend against such an attack than to launch one. Although a committed hacker may be capable of spending as much time as is needed to break into a system, a company cannot spend as much money as needed to defend itself against all possibilities.

The number of flaws that can be exploited is staggering, considering all the different platforms and devices that make up current information technology (IT) infrastructures. Compounding the problem is the fact that a hacker can work in relative anonymity using a $500 computer and modem to launch attacks. Even worse, hackers can work from any number of Internet kiosks available in airports, malls, cafés, and even laundries. As hackers get more sophisticated, and as more easily utilized scripts become available, attacks will be launched from mobile devices that can roam in and out of different geographic zones and then be discarded—making tracking of the attacker next to impossible.

21.4 WEB APPLICATION SYSTEM SECURITY. In spite of the fairly bleak picture painted here, organizations can effectively manage their risk from hackers. As in many other security domains, the security posture or stance assumed by the organization is critical for deterring and thwarting hackers. To use a physical-world analogy, consider burglars who intend to break into homes in a nice neighborhood. As the burglars scope out potential targets, they will notice some houses with burglar alarms—complete with conspicuous signs of the alarm systems—and some without. In all likelihood, the burglars will bypass the houses with the burglar alarms and move on to the other, less well-protected targets. Thus, the security stance assumed by the owner plays an important role.

Every organization must first determine its desired security stance as documented in its security policy. System administrators use the security policy to configure the systems, routers, firewalls, and remote access solutions. Without an explicit security policy, there is no way to determine what the security stance of the organization is, how to configure its systems, or even if a nonobvious security breach has occurred. Once the security policy is developed, the actual security implementation must be assessed. That is, the system must be tested and evaluated to determine how well it meets its security policy. Usually there is a difference between the desired stance and the actual stance. This difference is the security gap between where the organization would like to be (secure against particular threats) and where it actually is in practice.

The process of developing a security policy and evaluating the organization's systems against that policy will identify not only the gaps between the actual security stance and the desired posture but also weaknesses in the security policy itself. It is important to have an independent party, preferably an outside party, evaluate the security stance of the organization. A third party can fairly assess whether the organization's system upholds the security policy. If the group that develops the security policy or the system

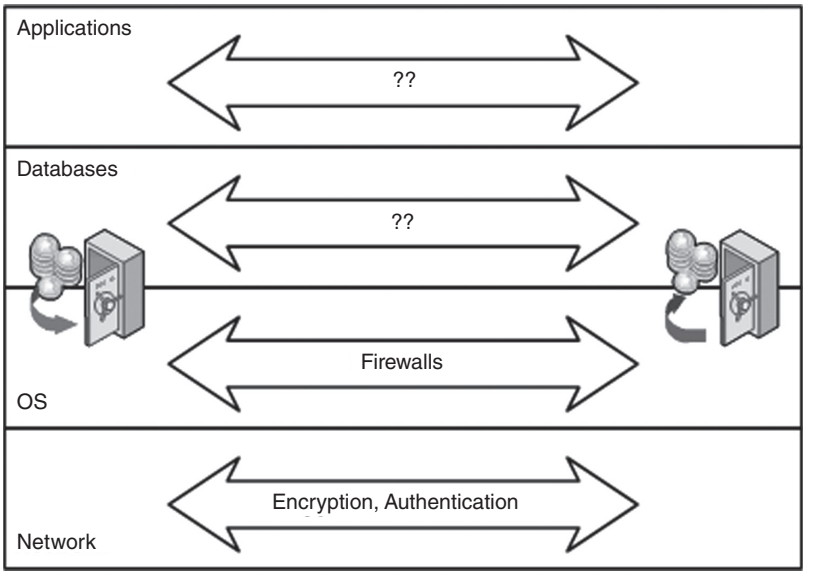


EXHIBIT 21.3 Layered View of an E-Business Application

configuration is also responsible for evaluating it, the evaluation may be biased, and potential vulnerabilities may be overlooked. This may occur not through ill will or dishonesty but rather because of the difficulty of being objective about one's own work.

21.5 PROTECTING WEB APPLICATIONS. Web application security can be understood from different views. First, consider the view of e-businesses in Exhibit 21.3. The diagram shows two e-businesses communicating over the Internet, perhaps performing business-to-business types of transactions.

In this view, the lowest layer of the e-business is the *networking layer*. At the networking layer, there are concerns about the reliability, integrity, and confidentiality of the data that runs over the communications channel. This layer is of particular concern because the Internet is a public switched network, meaning that any number of third parties may have access to the data that traverses the nodes of the Internet on the way from the data source to its destination. Also, the Internet Protocol (IP) is a connectionless protocol, which means there is no dedicated circuit between source and destination. As a result, packets sent during a given session may take different routes to their destination, depending on traffic congestion and routing algorithms. Because IP is an unreliable datagram protocol (i.e., IP uses independent packets forwarded from node to node, but there is no guarantee of successful transmission), the networking layer includes a connection-oriented reliable transmission layer such as Transmission Control Protocol (TCP) that ensures that dropped or lost packets are retransmitted and that bit flips that may have occurred during transmission (e.g., over wireless networks) are corrected.

Although TCP/IP provides for more reliable delivery of Internet packets, IP version 4 does not provide *secure* connection services. Typically, this means that there is no *guarantee* of confidentiality, identification, or even delivery of packets sent from one Internet host to another. Because packets often traverse several Internet nodes from source to destination, packet contents can be intercepted by third parties, copied, sub-

stituted, or even destroyed. This is the risk that most people citing risks of e-commerce have decried; they have overlooked the more substantive risks of e-commerce dealing with server- and client-side security and privacy. IPv6 holds promise to bring many security improvements to the network layer, but it has not been deployed widely in the United States. Fortunately, even in IPv4, we have good solutions to the data confidentiality problem. Cryptographic techniques can provide strong guarantees for data confidentiality, authentication of parties, and integrity of data sent during the transmission. Furthermore, digital signatures can be used to sign received mail in a "return receipt" application that provides guarantees of delivery of email. Thus, as shown in Exhibit 21.3, we can use encryption services to protect data transmitted over the network.

The operating system (OS), or *platform*, that hosts the Web applications lives on the networking layer. In a layered model, the services of one layer use the services of the lower layer and provide services to upper layers. The network layer often is thought of as a core portion of the operating system; however, from a layered services point of view, the OS software runs on top of the network layer.

Operating systems are notoriously rife with software flaws that affect system security. Operating systems are vulnerable because commercial OSs today are immensely complex; for instance, the Windows Vista operating system is purported to have more than 50 million lines of source code. It is impossible to catch all software design and programming errors that may have security consequences in a platform this complex. Even though UNIX operating systems have been in use for the better part of 30 years, new flaws in OS utilities are found on a weekly basis across all the different UNIX platform variants.[4]

Security holes in the platform are critical by nature. That is, if the OS itself is vulnerable to exploitation, security provided by the application can be compromised by holes in the platform. The OS is always the foundation on which applications are built, so cracks in the foundation make for weak security at the application layer. As Exhibit 21.3 suggests, firewalls provide protection against some operating system flaws. One of the key roles of firewalls is their ability to shut down services offered to *logical domain addresses*.

Using Internet domain addresses, the firewall administrator can partition Internet addresses into *trusted* and *untrusted* domain ranges. For instance, any Internet address outside the company's domain can be considered untrusted. As a result, all OS services, such as remote logins, can be shut down to everyone outside of the company's domain. Even within a company, the domains can be partitioned so that certain *subdomains* are trusted for access to certain machines, but others are not. The key benefit then of firewalls is their ability to restrict access to the platform through *offered services* (i.e., specific functions that pass data through the firewall). As a result, firewalls can make it easy to hide OS flaws from untrusted entities.

Even so, firewalls are vulnerable to data- or code-driven attacks through offered services. These are sometimes called "allowed path" vulnerabilities. For instance, an attack through SMTP (mail) or HTTP (Web) will not be stopped by a firewall if the firewall is configured to let email and Web services through, as is necessary for e-commerce. Firewalls also will not stop OS exploits from insiders or from the trusted entities that are granted access to the platform. However, it is important to realize that firewalls, if properly configured, can close down exposure to a significant number of platform vulnerabilities simply by denying untrusted outsiders access to platform utilities and services. Exhibit 21.3 illustrates reasonable protection from network- and platform-based attacks but not from application and database attacks. The database

layer is shown separately in the diagram because of the importance of its role in e-commerce; however, database attacks usually can be considered as a type of application attack.

Application-based attacks represent a critical issue that is gaining more attention in recent years. There is no simple solution, but there is a growing sense that the security provided by a firewall, an SSL-enabled Website, and encryption, together with digital certificates and signatures, is not enough. These measures are necessary but not sufficient. Applications, above all, *are* the online business. The Payment Card Industry Data Security Standards (PCI DSS) puts any organization processing credit card transactions on notice that Web application vulnerability protection was required as of June 2008. The PCI DSS specifies either "having all custom application code reviewed for common vulnerabilities by an organization that specializes in application security" or "installing an application layer firewall in front of Web-facing applications."[5]

Online applications are increasingly sophisticated, and the software that implements the application logic has become highly complex, requiring component-based and object-oriented paradigms such as Enterprise Java Beans, CORBA, and DCOM/COM services. Collectively, these are known as *application servers*. The key point, however, is that because the application logic is custom and complex, it is often rife with errors in implementation or logic that can be and often are exploited by hackers.

Application security must not be confused with marketing claims. A secure online application is one that is resistant to attacks. It is not simply one that authenticates end users, encrypts transaction data, provides nonrepudiation of transactions, and guarantees service. These are all matters of importance that address characteristics of the transaction, not properties of the software. The remainder of this chapter addresses the software problem in some detail. The next section provides a different view of e-commerce systems from the layered view discussed previously. It identifies vulnerabilities in the different software components and strategies for managing the risks.

21.6 COMPONENTS AND VULNERABILITIES IN E-COMMERCE SYSTEMS.

Exhibit 21.4 shows a generic multitier architecture of an e-business, together with a summary of the types of vulnerabilities and risks to each of the major components. Using the Internet, Web clients (running on PCs or handheld devices) interface with a front-end Web server, a middleware layer of business application logic, back-end databases, an ERP system, supply-chain management software, and even some legacy systems that are now brought to the Internet.

21.6.1 Client-Side Risks.

Most e-commerce is performed using standard Web browsers and mail clients. Increasingly, e-commerce is being performed on handheld mobile devices such as personal digital assistants (PDAs) and mobile phones. The security risks particular to wireless devices are covered in Chapter 33 of this *Handbook*. Client-side security risks are mainly from malicious mobile code such as Web scripts, ActiveX controls, and hostile Java applets.[6] Another major risk in client-side software is loss of privacy.[7] Each computer, with its related software, receives and transmits a great deal of personal identifying information (PII). For instance, browsers may convey information about the computers (name, IP address, browser type, version, company name) and sometimes about the users themselves, particularly if automatic form-filling features have been enabled. Browsers also are used to track movements through the Web. For instance, every Website the browser visits typically gets a record of the previous site from which the user entered. Banner ads in Web pages also track which sites have been visited in order to create a profile of Web usage, and cookies can be

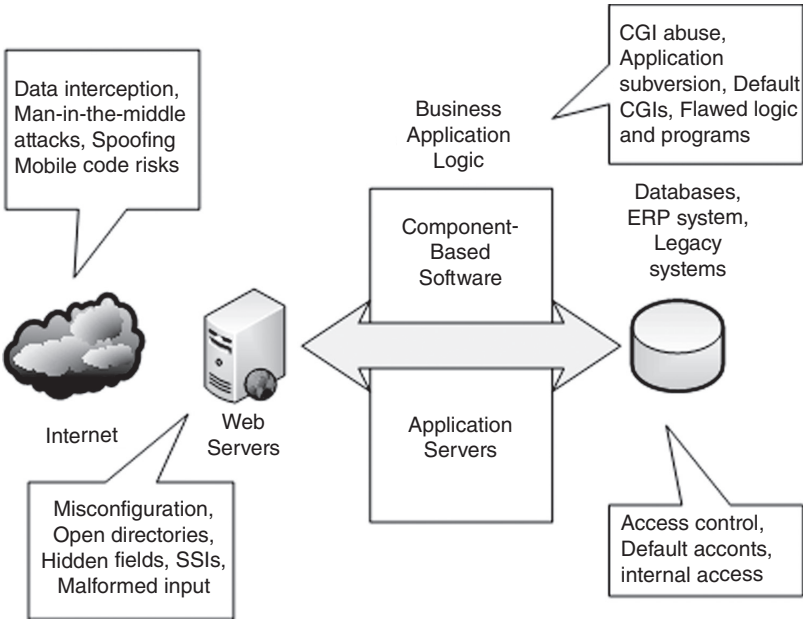


EXHIBIT 21.4 Multitier Architecture of an E-Business

placed on end-user systems for the purpose of tracking and at times re-creating user input.

A class of more insidious programs, known as *spyware*, can send information about computer usage out to specific sites, often without the user's knowledge or approval. One of the key risks with client-side software is that users simply do not know what the programs are revealing about themselves, or to whom. A simple principle is that a program should not be sending out any data that the user does not know about. An audit of the programs on any given machine probably would reveal that quite a few are in violation of that principle. Although most spyware programs are written to provide marketing data for software vendors or to redirect a user to predetermined sites, many manipulate and target users more effectively by spying on usage activity and even installing additional malicious code. For example, SpectorSoft markets a spyware tool known as eBlaster to people who suspect their spouses are engaging in illicit online affairs. This type of spyware program spies on user activity, from keystrokes to screen shots, and sends a log of the activity out over the network to a predetermined logging site. Spyware can be identified using diagnostic programs such as Ad-Aware from Lavasoft (www.lavasoft.com), but because of the dynamic nature of spyware technology, it may be difficult to identify all possibilities of infection.

A final client-side risk that businesses need to be especially concerned about is the risk of malicious executables that run on their user workstations. The desktop machine is like a petri dish for software: It is constantly changing and growing with new software executables—some of an unsavory nature. Malicious software, or *malware* as it is now known, finds many ways of infecting machines. For instance, one common way of disseminating malicious software is via email attachments. Another is by masquerading as legitimate software on a Web page available for download. Users often upload and download software to and from internal network file shares. In addition, USB "thumb drives" and even old-fashioned floppy disks are still a viable way of transmitting

malicious software. The Back Orifice "remote administration kit" is a well-known example of malicious software that, when installed, will allow anyone—including a hacker—to administer and control a machine remotely. Some malware products, running on an internal machine, may compromise the entire network because a remote hacker can control a trusted resource, possibly with enterprise administration rights. Data leakage and theft offered by the initial spyware infection can result in incidents that might undermine the integrity, authenticity, availability, or auditability of an entire data infrastructure. It is essential, therefore, that corporations carefully filter and closely monitor the application software that is downloaded and run (with or without the user's knowledge) on all systems from the e-commerce perimeter to the individual internal hosts.

21.6.2 Network Protocol Risks. Network risks primarily arise from sending confidential data over the Internet—a public, packet-switching network. Many good protocols address the risks of sending confidential data over the Internet.[8] In fact, a few years ago, the list included:

- SET
- SSL
- S/HTTP
- S/MIME
- CyberCash

Although some of these protocols are still around in one form or another, the industry has generally accepted SSL as the protocol of choice for secure Web browser transactions. The objective of most secure network protocols is to layer security properties on top of the TCP/IP network layers. Although TCP/IP provides reliable and robust delivery of datagrams over the Internet, it does not provide confidentiality, authentication, or strong message-integrity services. These are the properties that secure protocols provide. Some go even further. For instance, SET-compliant protocols leave the credit card number encrypted even at the merchant site. Since the merchant does not need to know the consumer's credit card number, by hiding the number from the merchant, a significant portion of credit card fraud can be eliminated. Rather than decrypting the credit card number at the merchant site, it is passed in encrypted form from the merchant to the credit-issuing bank. There it is decrypted, and the merchant's account is credited the amount of the purchase. The protocol details of SET, SSL, and other e-commerce protocols are described in Chapter 3 of this *Handbook*, in *E-Commerce Security: Weak Links, Best Defenses,* and in other books as well.[9]

Depending on the needs of an online business application, there is a requirement for more or less of the security properties afforded by secure protocols. For most Web browsing, a secure protocol is unnecessary; the standard Web protocol, HTTP, suffices. However, when customers send confidential or personal information to a site, a secure protocol that encrypts the data is preferable. The de facto secure protocol standard is SSL, now implemented in every standard Web browser. SSL will not only negotiate a secret session key between the Website and a client to encrypt the data, but it also will authenticate the Website. The Website must have a valid certificate endorsed by a Certificate Authority, which the users implicitly trust. Using a list of trusted Certificate Authorities maintained within the client browser, access can be prevented to other, untrusted sites. Once the connection is established, the user can verify that the Website

is, in fact, the one intended by examining the site certificate. Users rarely do this in practice, but the certificate is available and should be more widely utilized.

These secure properties—that is, encrypted sessions and host site authentication—serve the purpose for most online commerce applications. Some applications, though, may demand even stronger security services. For instance, banking and investment applications often transmit highly confidential information, with the possibility of huge financial losses through inadvertent error or intentional fraud. These types of transactions require not only confidentiality of the data, but also authentication of the client. A financial institution must never give access to account information or permit transactions without authenticating the user.

Common identification schemes on the Internet include simple user name and password authentication. A much more secure solution is to require strong client authentication using client certificates. SSL supports client certificates, although sites rarely use this capability because it involves requiring customers to obtain a certificate from a Certification Authority.

In the future, e-commerce protocols will need to be increasingly sophisticated if they are to meet the stringent security and privacy requirements of new Internet applications. For example, as criminal, medical, patent, and other important databases migrate to the Internet, the protocols developed for accessing them need to consider the security and privacy needs of the database owners and maintainers, as well as those of the client or requester. Today, the progress in genomics is producing much information about the likelihood of developing deadly disease based on a genetic DNA sequence. This knowledge raises moral and ethical questions as to how much information about potential diseases should be revealed to doctors and patients, and also raises the specter of such information getting into the wrong hands once it is accessible on the Internet.

Consider the case of a doctor querying an online genetic disease database with a patient's DNA sequence. The online application attempts to match the DNA sequence with diseases that might develop in the future. If the database were maintained by a commercial entity such as an insurance provider, the patient almost certainly would not want the company to know of any disease that might be returned as the result of a query, because that information could be used to deny both insurance and employment.

Likewise, the database maintainer probably would not want to know of the query or of its result, as such knowledge might put it at risk for lawsuits, should the information be leaked. Furthermore, the database maintainer would want the rest of the database to remain inaccessible except for specific results returned to an approved inquiry. Preventing access to any other information in the database would help to protect the commercial interests of the company, because then the database could not be duplicated easily. Nor could queries be made without a cost-tracking mechanism. To support this dual model of secure and private information access, e-commerce protocols need to be developed and commercialized that not only encrypt data in transmission but also consider the total security and privacy needs of both parties to the transaction.

Another e-commerce application area that will require better security and privacy protocols involves applications that accept e-cash or digital coin payments. Currently, most online payment schemes use either credit or debit cards, with payments made from the buyer's checking account at a bank. Payments are made either with online verification of funds or with off-line batch payments at the end of the day. A number of applications, particularly those involving payments of a few dollars or even pennies, are being created that cannot support the costs of a bank-based transaction.

Many commercial services and products, such as vending machines and parking meters, are coin activated and do not require customers to have an account or a line

of credit. Although efforts are being made to convert such services to computerized devices, with micropayment capabilities, it will be many years before this is actually accomplished. In any case, there will always be customers who want to pay with cash or its electronic equivalent, so that the transaction cannot be tracked by a third party, such as a bank or a creditor.

Newer online applications for micropayments may include collecting fees for downloading music, data, weather reports, stock quotations, articles from magazines, and pages from books. Many of these applications are provided today without charge and supported by banner advertising, but a concern for profits is motivating most Websites to seek additional sources of income. Whatever the application, there is a need for cash-based alternatives to the current account-based system for making payments. The key security and privacy concerns are with ensuring that e-cash is properly spent and accounted for and that anonymity is preserved throughout. Several protocols have been developed with these goals in mind, but none has reached commercial success or adoption by the vendor community. As mobile e-commerce begins to drive more traditionally cash-based transactions (e.g., parking meters, vending machines, and ticket booths), wireless vendors may adopt these new digital cash-based protocols.

Regardless of the network protocol used in an e-commerce application, the key concern is for those attackers who will attempt to breach the easiest obstacle in their quest to obtain system privileges and unauthorized access to data. If the security provided by the network protocol is perceived to be strong, attackers will look for alternatives that bypass the network security. For instance, these types of standard attacks from the hacker's toolkit will bypass the security provided by most e-commerce protocols:

- **Man-in-the-middle attacks.** Capturing transmissions in transit for eavesdropping or forgery
- **DNS attacks.** Altering records in the worldwide Domain Name System to misdirect connections to the wrong addresses
- **War dialing.** Automated testing of every phone number in a block of numbers
- **Exploiting software vulnerabilities in network services.** Such as FTP, Bind, SMTP, and HTTP servers
- **Internal access.** Improper use of authorized access by an insider (employee, contractor)
- **Leveraging trusted hosts.** Attacking from another system that has a privileged relationship with the target system
- **Brute-force crypto attacks.** Automated testing of all possible decryption keys to decipher a ciphertext

In summary, it is important not only to select the appropriate network protocol for each online application but also to avoid a false sense of security that might arise from use of a "secure" network protocol. Good security engineering will consider vulnerabilities in other components of the system that are more attractive targets to determined hackers.

21.6.3 Business Application Logic. The business application logic pictured in Exhibit 21.4 represents one of the key areas of vulnerability in e-commerce systems. The program logic encodes what the online business is offering in terms of products,

services, and customer convenience. It also defines the look and feel of the Website and provides all of the interactive features, such as dynamic Web pages, personalized Web pages, and online transaction capabilities. Because each application is unique, the software that implements the logic must be, in many cases, custom-developed for each particular site.

In contrast, most of the other software components of a Website are commercial off-the-shelf (COTS) software. For instance, the Web server, back-end databases, and supply-chain logistics software are often purchased off the shelf from software vendors. With COTS software, the end user has no control over the code and therefore is not responsible for coding bug fixes. When software bugs in COTS software are discovered, the software vendor usually issues a patch or incorporates a bug fix in the next release version. Software vendors can fix discovered bugs, but they depend on customer sites to actually apply the patches or upgrades to the software. In practice, this occurs less often than desired and is a significant reason why many Internet-based systems are vulnerable.[10] The most important task in securing COTS software systems is to make sure that: (1) they are properly configured for a secure installation according to the site's security policy, (2) the software is properly updated to the current version and patch level, and (3) administrators are aware of the risks associated with bugs in the software.

Because many business applications are custom-developed, either by an in-house staff or by outsourcing to an e-business developer, the programs represent a key risk for several reasons. The dynamic, interactive nature of e-businesses, coupled with increasingly sophisticated online services, requires a significant amount of development to code the logic. As a result, the application programs tend to be very complex pieces of software, likely to contain flaws, and susceptible to the kinds of attacks launched against Websites. In practice, errors in design and implementation of business application logic often compromise the security of an e-business.

Traditionally, the middle tier of software is implemented on Web servers using CGI and more recently PHP: Hypertext Processor (PHP). CGI and PHP scripts are programs that run on the Web server machine as separate processes from the Web server software. The Web server invokes these general-purpose programs in response to user requests. The CGI/PHP script's main function is to process user input and to perform some service, such as retrieving data or dynamically creating a Web page for the end user. Because CGI and PHP scripts process untrusted user input, the security risks associated with them and other forms of middle-tier software are extremely high. Many attacks against Web-based systems are implemented by exploited CGI/PHP scripts. And while CGI and PHP scripts can be written in any general-purpose programming language, they are written most often in Perl, C, Tcl, and Python.

More recently, component-based software (CBS) is making inroads in e-commerce applications as well as in standard Web applications. The purpose of CBS is to develop, purchase, and reuse proven software in order to implement application logic quickly, easily, and with high quality. Two of the more popular component frameworks for e-commerce applications are Enterprise JavaBeans (EJB) using XML and Java 2 Enterprise Edition (J2EE), which supports component-based Java. Other component models include the Object Management Group's (OMG) Common Object Request Broker Architecture (CORBA) and Microsoft's Common Object Model (COM) and Distributed COM (DCOM). These component frameworks are the glue that enables software components to use standard infrastructure services while hiding the details of the implementation by using well-defined interfaces.

Business application logic, when coded in CBS systems, usually runs on application servers with particular component models, such as EJB, CORBA, COM, and DCOM.

CBS also provides an interface to back-end services such as database management, enterprise resource planning (ERP), and legacy software systems.

In addition to supporting traditional CGI functions, component-based software is expected to enable distributed, business-to-business applications over the Internet. The component-based software paradigm also supports good software engineering, as described later. The Unified Modeling Language (UML) facilitates object-oriented analysis and design for component-based frameworks. In addition, as the market for component-based software expands, many standard business application components will be available for purchase off the shelf.

Although the benefits of component-based software are numerous, they pose security hazards similar to those of CGI scripts. Component-based software enables development in general-purpose programming languages such as Java, C, and C++, which can execute with all the rights and privileges of server processes. Like CGI, they process untrusted user input, and because component-based software can be used to build sophisticated, large-scale applications, the likelihood for errors may be even greater than for simple CGI scripts. Regardless of the implementation—CGI or application servers—the security risks of server-side software are great, and therefore server-side software must be designed and implemented carefully.

The key risks in the middleware layer of e-commerce sites are:

- Misconfiguration of the CGI
- Default and development CGI scripts being left on the production server
- CGI misuse
- Application subversion
- Flawed logic
- Programming errors

21.6.4 CGI Script Vulnerabilities. CGI scripts are frequent targets of attackers because they are often misconfigured and vulnerable to misuse.[11] When designing CGI scripts, it is prudent to expect the unexpected, particularly the malicious attack. Although the Web designer has control over the content of CGI scripts, there is no control over what end users are going to send to them. Also, often overlooked are vulnerabilities of CGI scripts that exist on the server as part of the distribution but that are not actually used in the application. Some CGI scripts, included as part of the Web server distribution, have well-known flaws that can be exploited to obtain unauthorized access to the server. Even if the default CGI scripts are not used as part of the Web server pages, anyone can access them by simply knowing the script names.

One of the most common—yet easily preventable—security hazards is misconfiguration of software, especially CGI scripts. One feature supported by many Web servers is the ability of individuals throughout an organization to write CGI scripts and have them execute from their own directories. Although useful for prettying up personal Web pages, this feature also can introduce system security hazards. In Web-based applications, the Web server should be configured to prevent CGI scripts from executing anywhere but in a single CGI directory under control of the system administrator.

The script-aliased CGI mode for Web servers ensures that CGI scripts will execute only from an explicitly named directory in the server configuration file. In addition, the CGI script path is not named in the Uniform Resource Locator (URL) to the CGI. Rather, the server "aliases" the explicit path to the CGI script to a chosen name, such

as *cgi-bin*. Thus, running the server in script-aliased CGI mode prevents rogue CGI scripts from executing while it also hides the explicit path to the CGI scripts.

The CGI script directories also should be properly configured using OS access controls. For instance, if CGI scripts are written in a compiled language such as C, the script sources should be excluded from the document root of the Web server, so that they cannot be accessed via the Web. They should be accessible to the system administrator or Web content development group only, and inaccessible to everyone else in the organization. If the script sources fall into the hands of malicious perpetrators, the source code can be inspected for flaws, making the perpetrator's job even easier. Access to the CGI executables directory, frequently called the cgi-bin, should be properly controlled as well. Only the Web server and administrator need access to this directory. Liberal access permissions to the CGI executables directory give those with malicious intent the opportunity to place their own scripts within the site.

Most CGI scripts are written in scripting languages such as Perl, JavaScript, and Python. Scripting languages are useful for rapidly prototyping systems, but they also let the developer write dangerous code very easily. For instance, it is easy to construct system commands with user input, a potentially dangerous situation. Writing the same system functionality requires several lines of programming code and knowledge of system libraries. The easy accessibility of scripting languages makes them appealing, but also threatening to security-critical applications. It is also important to prohibit access to interpreters from the Web server. For instance, system administrators may be tempted to include the Perl interpreter in CGI script directories; however, doing so provides direct Web access to interactively execute Perl commands—an extremely dangerous configuration.

Finally, administrators should account for every CGI program on the server in terms of its purpose, origin, and modifications. Remove CGI scripts that do not serve a business function, and view with suspicion CGI scripts that are distributed with operating systems and Web servers, downloaded from the Internet, or purchased commercially. These steps will eliminate most of the potentially dangerous CGI scripts. Once a stable set of CGI programs is established, make a digital hash of the program executables (e.g., using MD5 or SHA-1) to enable future integrity checks.

21.6.5 Application Subversion. Application subversion attacks are not discussed often in relation to e-businesses, but they represent a significant threat to most online applications. Application subversion is a form of program misuse. Unlike buffer overflow attacks, application subversion attacks exploit the program logic without violating program integrity, in order to elevate user privileges and gain unauthorized access to data. It is the very complexity of the target program that gives the attacker the means to gain unauthorized access. Application subversion attacks use programs in ways that the program's designers and developers did not anticipate. Typically, these attacks are not scripted, but rather developed from online interactive use and subsequent abuse.

Referring to Exhibit 21.1, an application subversion attack will attempt to discover ways of short-circuiting paths in the workflow. For instance, there may be a hidden path that lets the user gain access to account information without being authenticated to that account. Many such attacks work on the premise that access to confidential information is not properly authenticated.

Another common attack sends malformed input to a program. Many Web pages use forms extensively to drive the application, while the data input on the form is checked using client-side scripting. An attacker can take advantage of the fact that many online

application developers assume that the client is going to use the form properly and that the scripts will check all input sent to the site. The attacker can examine the data stream sent by the form and then, rather than using the form, send a modified data stream via URL within the browser or incorporating it into a custom command of the attacker's choosing. An attacker may be able to obtain access to the application by placing system commands in the input stream. If the input stream is subsequently used in a *system()* call by the online application, the end user may force the execution of system commands on the attacker's behalf.

Some application developers rely heavily on hidden fields in the HTML document. Hidden fields allow the Web developer to include information on the page that is not displayed, although the end user can see the hidden field data simply by viewing the HTML source. The mistake application developers make is, first, in believing that the end user cannot see the hidden fields, and, second, in relying on the integrity of the hidden field data for making online decisions. Some online merchants have made the mistake of including pricing information for items in the hidden fields and using those prices to determine the cost of the online transaction. The end user can simply change the pricing in the hidden fields and send lower prices back to the merchants, for a discounted purchase.

Another misuse of hidden fields is to redirect application output. For instance, some Websites include file system path information in hidden fields on their Web pages. This information is used by a server-side script to determine where to read or write transaction information. Attackers, simply by changing the hidden field, can overwrite files or read files to which they should not have access. In some cases, it may be possible to store a program entered in a form field to be used later as a means of running a privileged shell on the remote system.

In summary, rigorous software quality assurance is necessary throughout the design and development of all Web-enabled and e-business applications, including front-end Web pages, application middleware, and operating systems. Once the software is believed to be immune to application misuse and subversion attacks, the system administrator must perform other activities to ensure the security of the e-business middleware:

- All unnecessary scripts or application server programs must be eliminated from the production server.
- Source code of application middleware must be carefully guarded against download or unauthorized access.
- Proper configuration of the CGI and application middleware is necessary to ensure executable access only to the correct application middleware, with the lowest practical privilege level. Sanity checking of inputs to application middleware must be done to ensure that only well-formed input is accepted.
- Testing of application code with the use of programs such as WebInspect by SPI Dynamics (www.spidynamics.com) that are specifically designed to discover flaws within Web-based applications.

21.6.6 Web Server Exploits. Web server security has been written about and covered in detail, including in *E-Commerce Security: Weak Links, Best Defenses*[12] and *The Web Security Source Book*,[13] among other titles. Here we highlight some of the common exploits of Web servers used against many e-businesses.

21.6.6.1 Configuration. The key to Web server security is its configuration. Like other complex pieces of software, Web servers are highly configurable to meet the needs of any given site. By default, most software vendors configure the software application for maximum functionality but minimum security. Thus, by default, when the server is first started, it is likely to be more permissive than any given company's security policy would like. The principal premise for configurability is in the variations that exist among sites.

A correctly configured Web server is the result of a policy that defines what access is allowed, to which individuals, for each resource. This policy, in turn, is used to configure routers, firewalls, and all public servers, such as the Web server. Configuring the Web server, while necessary, is by no means sufficient to secure the system. The discussion of application server exploits in Section 21.6.5 demonstrates this principle.

21.6.6.2 HTML Coding and Server-Side Includes. Once the Web server is configured as securely as possible, it is important to ensure that the Web pages themselves do not open holes in the security. Many Web page developers fall into some common pitfalls that may compromise the site's security. The preceding section mentioned the problem of relying on hidden fields in HTML for security or business-critical data. Users can abuse the hidden field data to subvert an application.

HTML offers other potential vulnerabilities. One that is most often criticized is *Server Side Includes* (SSI). The SSI configuration option, if enabled, allows directives to be embedded in HTML that the server will execute. For instance, if the next statement were embedded in an HTML document, it would direct the server to display the contents of the system password file:

```
<!-#exec /bin/cat /etc/passwd ->
```

Certainly, Web pages should not be written with SSIs without a compelling reason. Although access to the HTML code is normally under control of the site, there are many ways an attacker might get an SSI into an HTML page. First, the attacker may have found another way into the system (e.g., by a CGI script exploit) but may want to provide either an easier backdoor or a redundant backdoor in case the CGI script vulnerability is found and closed. Once the CGI script is exploited, the attacker may implant an SSI directive within one of the site's HTML pages. Another way for an attacker to gain access is to have the server generate a Web page with the SSI of choice embedded in the HTML. How can an attacker do this? One approach exploits a server that generates dynamic HTML depending on the end user's data, preferences, or history. If the Web server ends up using some of the attacker's data to generate the HTML page, the attacker may be able to insert an SSI directive in the HTML. In summary, a better solution than continuously monitoring the HTML pages is simply to disable the SSI. In that event, even if SSIs were embedded, the server would not execute them. This is a configuration option, and like all system configuration files, the Web server configuration file should be protected by both file permission protection and file integrity checks, to ensure that it cannot be tampered with.

Although SSIs are often highlighted, a more common risk results from keeping documents or files in a publicly accessible portion of the Web server. The accessible portion of the Web server is called the *document root*. This root specifies the portion of the file system that the Web server can read and display to a Web client if requested. The document root can be a superset of the Web pages that are actually displayed when the user clicks through a Website. There may be other documents in the document root

that are not linked to or from Web pages. This does not mean they are not accessible, however. Simply giving the correct address for any document will result in either displaying or downloading the document to any Web client. Therein lies the problem.

It is worth noting that HTML content is no longer found only in files served by the Web. On the local hard disk, Microsoft Windows applications often install Help documents comprised of CHM ("chum") files and binaries. An attacker could exploit this vulnerability, leading to elevated privileges to the Windows XP "My computer" zone; unfortunately, local files are often overlooked and considered "trusted" during penetration tests.[14] Recognizing these risks, Microsoft Corp. began use of Microsoft Assistance Markup Language with the Windows Vista OS.

21.6.6.3 Private Documents in Public Directories. Private documents inadvertently placed in a public directory can result in a compromise of confidential information and loss of privacy. For example, if a database file of credit card numbers (say, *cardnumbers.mdb*) were stored in the document root of the Web server for a fictitious company, *mycompany.com*, this URL address typed in a Web browser could download the file: www.mycompany.com/cardnumbers.mdb

This risk is even greater if *directory browsing,* another configurable feature, is enabled. Directory browsing allows an end user to view the contents of the file system at a given directory level if a Web page does not exist with the same name. Directory browsing is really a way to explore the file system of another site using a Web browser. Users may view this feature if they go back one level from a given Web page by deleting the right-hand section of a page's URL and viewing its higher-level directory (e.g., in http://a.b.com/d/e.htm, one would remove the "e.htm" thus attempting to browse http://a.b.com/d/). Attackers can learn a lot of valuable information from viewing the contents of a directory including private files. Furthermore, the browser itself provides for clicking on a file name in the directory structure, which causes the file to be downloaded. Again, directory browsing, if enabled, is an open vulnerability and an unfortunately easy way to download private or confidential information.

21.6.6.4 Cookies and Other Client-Side Risks. Another potential vulnerability for an e-business is the use of *cookies*. Because HTTP is a *stateless protocol*, each new Web page that is visited has no memory of the last Web page that was visited by that user. Cookies are used to "keep state" between different Web pages visited in a given session. Cookies can make an e-business transaction appear to be seamless, sequential, and coordinated.

Most people, when discussing the risks of cookies, focus on the client-side privacy risks. Although these certainly exist, cookies also pose risks to the businesses that employ them. If the information contained in cookies is trusted, much the same way that the content in hidden fields is trusted, then the e-business may be vulnerable to cookie exploits called *cookie poisoning*. Some Websites use cookies to carry authentication information for a given user who traverses its pages. Once users have authenticated themselves, their token of authentication may be carried with them via cookies from one Web page to all subsequent pages at that site. Using cookies is a fairly weak form of authentication. The cookie can be stolen easily by someone snooping a user's local area network or the Internet and then, with the information gained, to access the user's personal pages on the Website. Secure protocols such as SSL should be employed to mitigate this risk.

Cookies also are used for other purposes that can introduce new vulnerabilities into critical transactions. Because cookies are under the control of end users, they can be

changed in whatever manner a user chooses. If cookies are designed to instruct a Web server where to write a customer-specific file, by changing the cookie data, an end user might overwrite other customer files or even replace critical system files. Similarly, if cookies are used for carrying order information, as is common in electronic shopping carts, then changing the contents of the cookies would corrupt the transaction. This could result in an unauthorized deep discount to the customer. This example and other cross-site scripting vulnerabilities are quite prevalent on today's Web. Browser exploits, HTTP header injection, cross-site request forgeries (CSRFs), and numerous flavors of phishing are all examples of Web client/user attack.

Regardless of the technology used, it is important to examine the features and functions of Web servers from a security-critical viewpoint. Dependence on a specific technology for critical transactions demands that the technology be trustworthy, so that it does not provide vulnerable points of attack.

21.6.6.5 Embedded Identifiers in URLs. Sometimes registration or unsubscribe messages include personalized URLs; for example, "To unsubscribe, click on http://www.some-company.com/unusbscribe/? = 12345" The problem is that the identifying code (12345 in the example) can be replaced by any other number—and some of the generated numbers will unsubscribe other subscribers. Another unfortunate use of such a technique occurs in email messages for participation in Webinars; sometimes the coded URL in the email goes directly to a registration page with information filled in—information such as the precise name, email address, and employer information drawn from the database of previous contacts.

In general, it is poor practice to provide such short identifiers that a user or a simple program or batch file could be used to delete valid records in a database or to harvest confidential data from a Website. For example, a simple html file in Adobe Acrobat can be used to force the program to access every URL in the file; if a malefactor has created, say, 20,000 unique identifiers and all or most of them are affiliated with real people, the malefactor could easily affect the accounts or collect the information shown on the customer-specific pages.

To avoid such problems, either create an address space with far more possible identifier values than a 1:1 mapping of real accounts to identifiers; e.g., for a 100,000-person list, use-generated values using a string of 20 positions with 94 values per position available in uppercase, lowercase, numbers, and special characters, resulting in $94^{20} = 10^{39}$ keyspace, far beyond anything manageable by an amateur.

21.6.7 Database Security. Databases traditionally have represented an important intellectual property of information-based companies. As a result, they have almost always been company proprietary and unavailable to public access. In the new model of business, however, many of these proprietary databases are made available on the Internet, often without careful consideration of the risks involved. The Web browser becomes the database query interface, often to unknown and untrusted entities.

Although there has been much research in database security over the last two decades, the commercial sector has adopted only two key tenets: authenticating and authorizing principals to certain objects. Access to databases thus is controlled by properly authenticating the credentials of the requesting principal and then verifying which objects the authenticated principal is authorized to access. Any online database application must perform these two functions rigorously to protect the most valuable assets of e-business as well as customer privacy.

Although many vendors claim secure channel access from the Web server to the database, there are many pitfalls. To start with, the fundamental reason why databases are vulnerable is that Web interfaces are commonly appended to what once may have been a closed and proprietary interface, without concern for overall security. Second, unsecured middleware programs such as CGI scripts or application servers usually mediate access from the Web server to the database. Web servers can provide client authentication from simple user name and password entry to strong certificated authentication. The needs of the business and the size of the accessing community will dictate which solution is feasible.

Despite obvious security advantages, most users do not store encrypted information in their databases, primarily for performance reasons. Encrypting and decrypting information on the fly during search, retrieve, and store operations can be too slow for real-time transactions. Also, even encrypting the data in a storage unit would not provide complete protection, as the online application must be able to read from and write to the database in clear text. Application-based attacks would still be able to get at the data while it was in plain, unencrypted text format.

Another key vulnerability of online databases also arises from application-based attacks. As described earlier, attacks that exploit vulnerabilities in the business-application logic often can provide unrestricted access to a database.

Attacks that exploit a buffer-overflow vulnerability in an application server program usually will be able to get command shell access on the remote server. From there, the attacker usually is able to find source code for the application server programs, such as Perl scripts or even C code, that are used to access the database. Because these programs *need* access to the database, they also must know the passwords used to access the various data partitions. If the programmers have foolishly *hard-coded* the passwords, then simply reviewing the source code may be enough to discover these passwords. With passwords in hand, the attacker can use the application server program via the Web interface to gain unauthorized access to the database. More directly, with a password it is possible to query a database from the command shell, using SQL commands or commands from the database language of choice.

Finally, like the other complex programs that run e-businesses, databases must be securely configured. Basic steps that the database administrator (DBA) needs to take include:

- Enforcing Web client authentication to the database
- Enforcing Web client authorization for access to database records
- Eliminating default database and database platform accounts
- Ensuring that passwords are read from encrypted files, not stored in program code
- Changing easily guessed passwords
- Configuring and maintaining internal access controls
- Auditing log files for suspicious activity

Account maintenance can be a key vulnerability in database management. Often the database software vendor will create a DBA account with an easily guessed password. Worse, DBAs may use the default account and password distributed with the database installation. This permits an attacker who has knowledge of the default passwords to gain access to all portions of the database by assuming the identity of the database administrator.

21.6.8 Platform Security. One area alluded to earlier in this chapter concerns the platforms that host components of an e-business. The platform, or operating system, represents the foundation of the e-business, but it is a potentially weak link in security. If there are cracks in the foundation, there is very little that even strong application software can do to keep the business secure. Therefore, it is imperative that system administrators properly patch platform vulnerabilities and maintain the security of the platform itself. As mentioned earlier, firewalls can go a long way toward blocking access to platform vulnerabilities by unauthorized outsiders. However, authorized but unprivileged users within the firewall can exploit known platform vulnerabilities to yield root privileges on a business-critical machine. Outsiders able to gain user privileges on the platform through any means also may be able to penetrate platform holes into severe security breaches.

Some key steps necessary to maintain platform security include:

- Eliminating default accounts generally installed with the operating system
- Prohibiting easily guessed passwords
- Enforcing password expiration
- Deactivating any unnecessary services that may be running by default
- Regularly applying security patches to the operating system
- Updating the operating system to its most recent release
- Ensuring that file access permissions are properly enforced, so as to prevent unnecessary access to critical files
- Enabling audit logging with intrusion monitoring
- Running system file integrity checks regularly

Some administrators believe that deploying a firewall is an acceptable substitute for configuring their platforms securely. Furthermore, with the plethora of different platforms running their enterprises, many system administrators give up on installing the latest OS patches and on securely configuring their platforms. They mistakenly assume that firewalls will protect them from all threats, even without maintenance. Unfortunately, relaxing host security can make the job of hackers easy. By hopping from machine to machine, they usually can find the valuable information they are looking for, or if that is their goal, they can wreak maximum damage.

An increasingly important platform for considerations of Website security is smart phones. By the end of 2012, there were 2.1 billion mobile Web users in the world at the end of 2012.

According to estimates by The ITU (June 2012), there were 2.1 billion active mobile-broadband subscriptions in the world. That is 29.5 percent of the global population.

- Mobile-broadband subscriptions have grown 40 percent annually over the last three years.
- Mobile-broadband subscription outnumber fixed broadband subscriptions 3:1.
- In developed countries mobile-broadband users often also have access to a fixed-broadband connection, but in developing countries mobile broadband is often the only access method available to people.

- Ericsson (November 2012) forecasts that global mobile broadband subscriptions will reach 1.5 billion at the end of 2012, and 6.5 billion in 2018. The mobile phone will continue to be the dominant mobile broadband access device.[15]

Tablets also have their own versions of browsers. Web designers are increasingly having to cope with significant differences among platforms accessing their code.

Craig Smith included the following points in a January 2013 article entitled "Optimizing Ecommerce for Tablets and Smartphones":

- **Improving navigation and usability**: Users should be able to select options easily and correctly; avoid drop-down menus and buttons that are close together.
- **Leveraging responsive design**: Determine how to present the Website according to what kind of device is accessing it.
- **Determining the purpose of the access**: Distinguish between the types of queries that are most common—and different—across platforms. Optimize design for the most common types for each platform.[16]

Another question is whether e-commerce sites should depend on Web access or develop applications (*apps*) specifically for smartphone and tablet operating systems.[17] A major advantage of apps is that they can be programmed to avoid the vulnerabilities common to mobile devices in these early years of widespread adoption.

21.7 SUMMARY. This chapter has provided an overview of the weakest links in Web applications, including Web clients, network protocols, front-end Web servers, back-end databases, application servers, and the platforms on which they run. Secure network protocols are necessary but certainly not sufficient for securing e-commerce. The vulnerabilities described here are largely based on software flaws that exist in the application layer of e-commerce transactions.

Perhaps the most common vulnerability in e-commerce systems is misconfiguration of software. Because the responsibility for software configuration lies with the user, a security policy must be implemented and enforced. Once a system is configured, it is important to subject it to third-party validation and testing. A third-party audit can ensure that the configured system, including routers, firewalls, servers, and databases, meets the specifications of the security policy. The system also should be tested periodically against well-known, common attacks as well as against newer threats as they arise.

Like any software system, commercial off-the-shelf software has flaws, many of which are security-critical. It is imperative that system administrators stay current with vendor Websites, with hacker sites, with newsgroups, and with wherever threats and patches to their software are released. Both the security and hacker communities are constantly at work finding flaws in COTS that software vendors and security firms generally correct quickly. Software consumers—those who buy and use the software—must do their part by applying all relevant patches to their vulnerable software.

For custom-developed software, which includes front-end Web pages, application servers, CGI scripts, and mobile content such as ActiveX or Java, developers must do everything possible to ensure that their software is not vulnerable to attack and that it does not infringe on users' privacy.

21.8 FURTHER READING

Anderson, R. *Security Engineering: A Guide to Building Dependable Distributed Systems*, 2nd ed. Wiley, 2008.

Bhargav, A. and B. V. Kumar. *Secure Java: For Web Application Development.* CRC Press, 2010.

Cobb, Stephen. *Privacy for Business: Web Sites and Email*. Dreva Hill, 2002.

Elliott, E. *Programming JavaScript Applications: Robust Web Architecture With Node, HTML5, and Modern JS Libraries*. O'Reilly Media, 2013.

Gabarro, S. A. *Web Application Design and Implementation: Apache 2, PHP5, MySQL, JavaScript, and Linux/UNIX*, 2nd ed. Wiley, 2013.

Gookin, D. *Android Phones for Dummies*. For Dummies, 2013.

Gunasekera, S. *Android Apps Security*. Apress, 2012.

Harwani, B. M. *Android Programming Unleashed*. Sams, 2012.

Kissoon, T. *Securing Web Applications*. Auerbach, 2013.

Long, F., D. Mohindra, R. C. Seacord, D. F. Sutherland, and D. Svoboda. *The CERT Oracle Secure Coding Standard for Java*. Addison-Wesley Professional, 2011.

Mark, D., A. Horovitz, K. Kim, and J. LaMarche. *More iOS 6 Development: Further Explorations of the iOS SDK*. Apress, 2012.

McGraw, G., and E. Felten. *Securing Java: Getting Down to Business with Mobile Code*, 2nd ed. Wiley, 1999.

Musciano, C., B. Kennedy, and E. Weyl. *HTML5: The Definitive Guide*, 7th ed. O'Reilly Media, 2014.

Scambray, J., Shema, M., and C. Sima. *Hacking Exposed Web Applications*, 2nd ed. San Francisco, CA: McGraw-Hill Osborne Media, 2006.

Six, J. *Application Security for the Android Platform: Processes, Permissions, and Other Safeguards*. O'Reilly Media, 2011.

Sullivan, B. *Web Application Security, A Beginner's Guide.* McGraw-Hill Osborne Media, 2011.

Welling, L., and L. Thomson. *PHP and MySQL Web Development*, 5th ed. Addison-Wesley Professional, 2013.

Zalewski, M. *The Tangled Web: A Guide to Securing Modern Web Applications.* No Starch Press, 2011.

21.9 NOTES

1. This original version of this chapter, entitled "E-Commerce Vulnerabilities," was adapted from Anup K. Ghosh, *Security and Privacy for E-Business* (Hoboken, NJ: John Wiley & Sons, 2001), chapter 4; adapted by permission.
2. Anup K. Ghosh, *E-Commerce Security: Weak Links, Best Defenses* (New York: John Wiley & Sons, 1998).
3. C. Cowan, P. Wagle, C. Pu, S. Beattie, and J. Walpole, "Buffer Overflows: Attacks and Defenses for the Vulnerability of the Decade," paper presented at DISCEX 2000, January 25–27, 2000, Hilton Head, S.C. Proceedings of the DARPA Information Survivability Conference and Exposition (Los Alamitos, CA: IEEE Computer Society Press, 2000).
4. Cowan, et al. "Buffer Overflows."
5. PCI Security Standards Council, Payment Card Industry Data Security Standard, v1.1. (September 2006), retrieved from www.pcisecuritystandards.org/pdfs/pci_dss_v1-1.pdf

6. G. McGraw and G. Morrisett, "Attacking Malicious Code: A Report to the Infosec Research Council," IEEE Software 17, No. 5 (September/October 2000): 33–41.
7. Ghosh, *Security and Privacy for E-Business*, chapter 7.
8. B. Schneier, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, 2nd ed. (New York: John Wiley & Sons, 1995).
9. Ghosh, *E-Commerce Security*.
10. W. A. Arbaugh, W. L. Fithen, and J. McHugh, "Windows of Vulnerability: A Case Study," *IEEE Computer* 33, No. 12 (December 2000): 52–59.
11. L. Stein, *Web Security: A Step-by-Step Reference Guide* (Reading, MA: Addison-Wesley, 1998).
12. Ghosh, *Security and Privacy for E-Business*.
13. A. Rubin, D. Geer, and M. Ranum, *The Web Security Sourcebook* (New York: John Wiley & Sons, 1997).
14. Michael Howard and David Leblanc. *Writing Secure Code*, 2nd ed. (Redmond, WA: Microsoft Press, 2002).
15. mobiThinking, "Global Mobile Statistics 2013 Part B: Mobile Web; Mobile Broadband Penetration; 3G/4G Subscribers and Networks," http://mobithinking.com/mobile-marketing-tools/latest-mobile-stats/b#mobilebroadband
16. C. Smith, "Optimizing Ecommerce for Tablets and Smartphones," *Practical ecommerce*, January 15, 2013, www.practicalecommerce.com/articles/3869-Optimizing-Ecommerce-for-Tablets-and-Smartphones
17. D. Traxler, "Mobile Commerce: Website or App?" *Practical ecommerce*, January 1, 2013, www.practicalecommerce.com/articles/3862-Mobile-Commerce-Website-or-App-

CHAPTER 22

PHYSICAL THREATS TO THE INFORMATION INFRASTRUCTURE

Franklin Platt

22.1 INTRODUCTION 22·2

22.2 BACKGROUND AND PERSPECTIVE 22·3

22.2.1 Today's Risks Are Greater 22·3
22.2.2 Likely Targets 22·4
22.2.3 Productivity Issues 22·4
22.2.4 Terrorism and Violence Are Now Serious Threats 22·4
22.2.5 Costs of Damaged IS Infrastructure 22·5
22.2.6 Who Must Be Involved 22·5
22.2.7 Liability Issues 22·6
22.2.8 Definitions and Terms 22·7
22.2.9 Uniform, Comprehensive Planning Process 22·8

22.3 THREAT ASSESSMENT PROCESS 22·9

22.3.1 Set Up a Steering Committee 22·9
22.3.2 Identify All Possible Threats 22·10
22.3.3 Sources of Information and Assistance 22·11
22.3.4 Determine the Likelihood of Each Threat 22·12
22.3.5 Approximate the Impact Costs 22·12
22.3.6 Costs of Cascading Events 22·13
22.3.7 Determine the Vulnerability to Each Threat 22·13
22.3.8 Completing the Threat Assessment Report 22·14

22.4 GENERAL THREATS 22·14

22.4.1 Natural Hazards 22·15
22.4.2 Other Natural Hazards 22·16
22.4.3 Health Threats 22·16
22.4.4 Man-Made Threats 22·16
22.4.5 Wiretaps 22·18
22.4.6 High-Energy Radio-Frequency Threats 22·19

22.5 WORKPLACE VIOLENCE AND TERRORISM 22·21

22.6 OTHER THREAT SITUATIONS 22·22

22.6.1 Leaks, Temperature, and Humidity 22·22
22.6.2 Off-Hour Visitors 22·22
22.6.3 Cleaning and Maintenance Threats 22·23
22.6.4 Storage-Room Threats 22·23
22.6.5 Medical Emergencies 22·23
22.6.6 Illicit Workstation 22·24
22.6.7 Other Local Threats 22·24

22.7 CONFIDENTIAL THREAT INFORMATION 22·25

22.8 SUMMARY 22·26

22.9 FURTHER READING 22·27

22.10 NOTES 22·27

22.1 INTRODUCTION. This chapter describes the wide array of possible physical threats that can affect information systems (IS) infrastructure. A *threat* is any credible situation—whether actual, imminent, predicted, or possible—with the potential for causing harm, damage, or disruption. The terms *threat* and *hazard* are used synonymously in this chapter and have the same meaning. *Risk* refers to the probability of occurrence of a threat. For a structured terminology for discussing threats and risks, see Chapter 8 in this *Handbook*.

The infrastructure affected can be any component of a computer system or communications network; any of the cables, wiring, or devices that transfer power or data; or any of the support services or utilities needed to sustain full IS performance. In addition, this chapter considers threats against employees as a significant component of physical security.

The speed and accuracy of any system is dependent on the performance of a long chain of physical components as well as on the productivity of all of the people who utilize or maintain each component. Anything less than full system-wide performance can be costly.

A physical threat is any event that can degrade the performance of an information system—whether such an event is actually occurring or imminent, or credibly likely or possible, or completely unexpected and without warning. The list of possible threats is a long one, beginning with natural and man-made external events that can degrade IS performance. Many other possible threats are internal, resulting from accidents or misuse or deliberate attack. Internal threats also include reliability failures, installation or maintenance issues, or lack of proper testing. Threats can also be caused by situations within the facility, building, or complex, or as a consequence of local or regional events, or, increasingly, events worldwide and even from outer space. Threats happen and can come from almost anywhere, so it is prudent and far less costly to try to anticipate all likely and significant threats.

Although in this chapter the terminology often refers to "all possible threats," readers must understand the implicit assumption that they are considering only reasonable categories of threat. Planning to cope with the results of a world-destroying impact with a continent-sized asteroid is useless; trying to define a strategy for responding to an attack by extraterrestrial invaders armed with ray-guns is impossible and useless.

This chapter begins a threat-assessment process, which starts by identifying all reasonably defined threat situations. This process must be comprehensive and rigorous and involve all stakeholders. It is not enough simply to assign likelihood and an impact for the usual or the obvious threat situations. There must be a full risk analysis, so that risk, vulnerability, and response and recovery costs can be quantified, as explained in Chapter 23 in this *Handbook*. Otherwise, the whole process is valueless. Absent comprehensive planning, security becomes an irrational selection of vendors and solutions, which can only add cost and increase liability risks. Alternatively, a strategic risk-management process can maximize future profits and add value, protect morale and productivity, enhance goodwill, and improve relations with customers and communities.

The events of September 11, 2001, sounded a shattering wake-up call that unexpected threats can actually happen and that even the best security practices, products, and services are of little value without proper planning, implementation, and support. Hurricane Katrina, four years later in 2005, demonstrated again that both business and government are still unprepared, even though such an event was predicted many times.

This chapter suggests a comprehensive perspective on physical security that can add value and help avert disaster. Chapter 23 in this *Handbook* goes on to suggest physical

security implementation and management that can protect people and property while optimizing morale, productivity, and profitability.

22.2 BACKGROUND AND PERSPECTIVE. Historic data and statistics are of limited value in predicting future threats. For a discussion of the unreliability of computer-crime statistics, see Chapter 10 in this *Handbook*. The past is no longer prologue because many new threats are emerging and many types of incidents are increasingly severe, widespread, complex, damaging, and costly. Proliferating and increasingly dispersed infrastructure system components are often vulnerable and hard to protect. And hybrid configurations of new and legacy systems vastly complicate the protection process.

There are few threat statistics that can reliably predict the infrastructure future. Computer-crime reports mostly cover logical security, while general crime statistics often do not relate directly to computer security. Historic information is further flawed because many incidents are never detected and far more are not reported for fear of embarrassment, liability, or loss of business. Many security incidents are masked as quality-control problems for the same reasons, or are misdiagnosed because no one had time to determine the true cause(s). Lacking reliable precedents, predicting future threats is especially difficult—yet increasingly necessary. Chapter 10 discusses computer crime statistics in more detail; Chapter 58 and Chapter 62 discuss risk-management methodologies.

Most incidents happen suddenly, without warning, and often where least expected. Many threats once thought to be unlikely now occur widely and strike with surprising intensity and devastation. Other contributing factors include poor risk management due to inexperience, denial, or complacency that can quickly turn routine threats into costly incidents. Businesses with good security preparedness can usually survive, while many others will not.

22.2.1 Today's Risks Are Greater. Today's threats are increasingly sophisticated, unpredictable, potentially serious, and increasingly commonplace as well. Disruptive incidents can result from mistakes or accidents, snoopers or hacking, vandalism, disgruntled or disruptive persons, labor disputes, demonstrations, civil unrest, extremists of many stripes, and, increasingly, both domestic and international terrorism. Although violent-crime incidence has decreased in recent years in some parts of the world, these data may be misleading because they rarely include workplace-related events. Violence in the workplace is now common and often without warning. Incidents can include harassment, bomb scares, robbery, hostage situations, shooting, or arson. And each and every one of these threats can seriously affect IS performance.

Physical threats can extend beyond direct attacks. Many scares do not result in actual violence, and disruptive incidents can occur outside of the workplace. However and wherever they happen, violence-related incidents are increasingly commonplace, disruptive, and costly.[1]

Those who threaten the IS infrastructure must at least gain access to some physical part of it, but often this can be done inconspicuously, by trickery, deception, or simply forced entry. A physical attack may be the best way to compromise an information system—much more effective and less likely to be detected than a logical attack. Often, many vital system components are vulnerable, exposed, or easily accessible. These components include wiring and cable runs, connection and junction points, system and network equipment, and the utilities that support them. Attacks or spying by physical means are often easy, fast, safe, and sure.

22.2.2 Likely Targets. Businesses and organizations are increasingly likely to be the targets of hackers, disgruntled employees, competitors, disturbed persons, demonstrators, hate groups, extremists, and even terrorists. Motives may include a conspicuous, tempting, or challenging target, rage or revenge, an opportunity to cause reputational damage or at least adverse publicity, to make a political statement, for extortion or blackmail, or for personal gain or profit. Often, there are no discernable motives. And beyond being likely targets, most businesses are convenient, easy, and safe targets as well, because most are unaware and unprepared for today's threats, much less for future ones. Although government facilities remain preferred targets, many are now better protected than most businesses, thanks to new procedures that are covered in Chapter 23 in this *Handbook*.

Another likely threat arises from the need of extremists and terrorists to finance their activities. Many groups and all independent, self-directed cells are dependent on crime to finance their operations. Today's crimes can include robbery and holdups, counterfeit currency and credit cards, Internet phishing and scams, theft, extortion, blackmail, and selling pirated software and other knock-offs. In addition, many foreign governments, businesses, and criminal organizations are actively engaged in spying. Although these are mainly corporate security and logical security problems, they also represent potential physical threats that must be deterred.

22.2.3 Productivity Issues. Good security can be directly correlated with high productivity, which, in turn, can improve performance, customer satisfaction, and goodwill. Good security is measurable and can strategically enhance each of these factors to add both value and profit. Anything less than good security invites wasted time and money.

People who do not feel safe will not be productive. This applies to employees, visitors, vendors, and others on premises, as well as to customers, vendors, stockholders, and other stakeholders at remote locations. Everyone using any information system must be comfortable that physical safety and privacy are insured and that the system is uninterruptible, secure, and operating at full performance. Everyone concerned must be involved in the planning process, generally understand the threats, and support the security procedures. Otherwise, performance inevitably suffers.

Whenever a security incident occurs, morale and productivity are likely to plummet and can remain low for many weeks, or even months. Whether the infrastructure is actually affected or not, significant disruption of operations is likely. Even when there is no injury or damage, the perception of a potential event can be costly; it can disrupt productivity, lose business and customers, and jeopardize goodwill.

No one knows how often productivity-related events occur. Businesses generally do not report them, and neither do the media. But these things do happen and probably with considerable frequency. Yet all threats can be mitigated to some extent and many more prevented at far less cost than the consequences by effective technical preparations, adequate employee awareness and training, and well-planned and well-rehearsed responses.

22.2.4 Terrorism and Violence Are Now Serious Threats. Acts of terrorism and violence are now a reality that can occur anywhere in the world. September 11, 2001, and the events that followed have brought home the stark reality that violence can happen anywhere and can cause massive damage and disruption. And most major disasters can disrupt information systems far removed from the actual incidents.

Workplace violence is also happening with increasing frequency and often at facilities assumed to be safe.[2] Bomb and biological or chemical scares, personal threats, harassment, hostage situations, and shootings are all happening with increasing frequency. Whether actual violence occurs or not, the threats alone, the many rumors generated, and the imagined proximity to danger are all productivity-related events that can seriously disrupt the performance of information systems for a long time. Therefore, these become infrastructure security issues that require special planning and should not be left to premises security personnel to prevent.

There are other serious threats from foreign intelligence, terrorists, and domestic groups generally unknown to the public. Some of these are explained well in the Project Megiddo report published by the Federal Bureau of Investigation (FBI) in 1999 in anticipation of the millennium. The report provides "an FBI strategic assessment of the potential for domestic terrorism in the United States undertaken in anticipation of or response to the arrival of the new millennium."[3] The threats cited then are basically unchanged today, except that more previously unknown threats have since been added. There are consultants with FBI contacts who can offer valuable advice.

Attempted violence is now a serious threat to all IS infrastructures. However, thorough security planning can do much to avoid trouble and needless expenses.

22.2.5 Costs of Damaged IS Infrastructure. Direct costs of system downtime can exceed many thousands of dollars per hour. The losses include the slack-time costs of people who cannot use the systems, costs of support and maintenance people diverted to restoring operations, recovery expenses, overtime, and often lodging, food, and travel expenses during recovery. Usually, many outside resources are needed for response and recovery. Everything quickly becomes expensive. Often, to further compound the costs, needed resources are just not available immediately.

Indirect costs can also be significant. Reestablishing and keeping good public relations can be expensive. Often, public announcements, news releases, and briefings to the news and financial media and to stockholders are needed to neutralize public embarrassment and control rumors. Key customers must be contacted and reassured, and pending orders rescheduled. Still more costs include lost business or market share and dropping stock prices. Competitors often will take as much advantage as they can, which necessitates further costs defending brands and reputation.

Any number of such costs can devastate an enterprise unless strong security measures are deployed effectively and quickly. Ad-libbed responses are often disastrous. In reality, an infrastructure outage of more than a few hours is often fatal, and the enterprise can never fully recover. Good security usually can prevent disastrous costs.

22.2.6 Who Must Be Involved. Many threats to information systems also involve corporate or premises security personnel, whose role is to protect people and property within and surrounding the workplace. The CEO and the CFO are also involved, because they must now comply with the applicable laws and regulations. Compliance requires that events that might materially affect financial performance must be included within an organization's financial filings and statements. Major security incidents are clearly such events whose impact must now be predicted—as described in this chapter.

Premises security often includes little more than guards, access control, and some surveillance. Their understanding of the IS infrastructure's special security needs is often minimal, lumped in with overall physical security procedures, and rarely extends much beyond the immediate premises. Yet, IS security requires additional knowledge,

experience, protection, and support. Good security must be strong, fast-acting, focused on specific targets, and closely monitored. Effective early warning systems are necessary in order to prevent threats from happening. In reality, each security function will have its own needs and priorities and use its own resources. During a serious incident, security for the premises, occupants, information systems, and the infrastructure must all coordinate efficiently and effectively. They must also work smoothly with local fire and police departments, with other emergency responders, and with many outside resources.

Who, then, should manage the process of determining the threats to the IS infrastructure? And who are the stakeholders who should be involved in this process?

The best person to manage the physical security of information systems is one who knows a great deal about possible threats and about the IS infrastructure. The office manager, facilities manager, or corporate security director or their staffs are usually ill-equipped to determine or manage IS security. Often, too, the chief information officer (CIO), chief information security officer (CISO), and IS security officers (ISSOs) deal with protecting data and data processing, and are not the best persons to understand IS physical security, especially in a large installation. Therefore, the best person is a trusted individual with the right knowledge and experience and with enough time to manage the process well. Planning and managing infrastructure security, implementing and testing it, training and security awareness, monitoring, and periodic updating of the system and procedures are a full-time job in most organizations, and require a support staff in larger organizations.

Chapter 65in this *Handbook* discusses the role of the CISO in detail.

Another consideration is that no one person should know all the secrets, which is a rule that has become especially important in dealing with the IS infrastructure. When trouble comes, many experts must mobilize very quickly, efficiently, and effectively. To do this, the responders must be familiar with all the information systems and have fast access to the infrastructure. No matter how well trusted, no one person should know everything about the physical and logical defenses. (If need be, such information should be safely stored with strong access controls.) It is wise to divide the secrets so that no one group knows or has access to them all. Having done this, multiple persons can then share each portion of the secrets and observe each other, so that no one person is indispensable. Chapter 45 in this *Handbook* presents details of employee management and security.

But only the security vulnerabilities and defensive implementations should be secrets. There is no point is letting others know about the defenses and where the organization is vulnerable. However, the process of determining the underlying threats should be common knowledge among all the stakeholders involved. By identifying a large number of threats that have been assessed—but keeping the likelihood, vulnerability, and impact information secret—we can discourage would-be troublemakers by at least providing some idea of the *number* of threats that have been considered. With any luck, they will attack less-prepared sites.

22.2.7 Liability Issues. Aside from the need for efficient emergency response, another issue that is becoming increasingly important, is potentially very costly, and is often overlooked. This is the issue of liability. Every organization has a legal and fiduciary duty to protect the people and property within and surrounding its premises. If any injury or damage occurs in or around a workplace—even long afterward—allegations of negligence will likely follow. The motive is often that damages awarded by the courts can be very large and lawyers often take on these cases on speculation, in hopes of

receiving very large fees. And whether the organization was actually at fault or not, the resultant legal fees, the time and costs needed to defend the organization, bad publicity and possible loss of business, and the eventual fines and awards can be devastating.

If negligence is alleged, the issue is whether the organization was properly prepared for the emergency and whether its response was effective. The question is simply: Did management perform its duty to protect the organization? An affirmative answer would require at least:

- Evidence of a thorough threat assessment process
- Good security plans, policies, and procedures that have actually been implemented
- Proper training and current security awareness
- Periodic drills, exercises, security reviews, and feedback from known events
- Periodic updates to assure that security remains effective

If it can be demonstrated that all these measures were not taken, or that there were deviations from generally accepted standards, the result could be punitive as well as compensatory damage awards. Gross negligence will probably be alleged as well, in which case insurance may not defend the accused organization or individuals, who would then be personally liable. Even with insurance in effect, it may not be sufficient to cover the very large penalties that juries commonly award.

Chapter 60 in this *Handbook* discusses insurance for responding to damages in IS.

Even when all the correct answers are given, the accused often find they are considered to be guilty until they can prove themselves innocent. And this can be a long, painful, and costly process.

However, a nearly iron-clad defense against liability is that federal procedures were followed. The most comprehensive of these are from the Department of Homeland Security (DHS) and its Federal Emergency Management Agency (FEMA) Directorate.[4] Of all of the other security planning and management procedures, which are many and varied, only the DHS/FEMA methodology is likely to become generally accepted. The procedures are already well known, uniform, and comprehensive, and they consider all threat possibilities and all response resources. It is unlikely that a plaintiff's attorney will ever allege that these procedures are deficient or ever bring a case once it appears that an organization is compliant. This is not to remove all liability, merely to reduce the scope to situations usually covered by insurance. Chapter 23 in this *Handbook* outlines how to plan, implement, and manage the DHS/FEMA procedures.

22.2.8 Definitions and Terms. Information infrastructure security is simply a means of IS performance assurance. Good security precludes any IS disruption that might degrade performance in any way. Any slowdown or loss of productivity, loss of data, breach of privacy, or disruption of the systems, networks, and utilities that support any information system diminishes performance. Good security assures that all systems remain fully operational, robust, and accurate, and that all their data remain private and cannot be compromised.

There are three elements of information systems security. Each element must be especially designed and maintained to protect against different threats of varying scope and intensity. The three elements include:

1. **Logical security**, which is also known as information systems security, protects only the integrity of data and of information processing.

2. **Physical security**, which is also called infrastructure security, protects the rest of the information systems and all of the people who use, operate, and maintain the systems. Physical security also must prevent any type of physical access or intrusion that can compromise logical security.
3. **Premises security**, which is also known as corporate or facilities security, protects the people and property within an entire area, facility, or building(s), and is usually required by codes, regulations, and fiduciary obligations. Premises security protects broad areas. It often provides perimeter security, access control, smoke and fire detection, fire suppression, some environmental protection, and usually surveillance systems, alarms, watchmen, and guards. Premises security is often an extension of law enforcement.

Clearly, there is much overlap between each element, and a threat to one is often a threat to the others as well. However, each element is likely to perceive and handle each threat differently. Although the remainder of this chapter deals with physical security, some threats that are usually considered the realm of logical or premises security must also be covered.

22.2.9 Uniform, Comprehensive Planning Process. The threat assessment process can go by a number of names, which are all functionally similar. Among the terms are emergency, disaster, operations, contingency and crisis-response planning, and damage control. Some processes are proprietary and do not share a common language or standardized procedures that can be understood by everyone involved. Many regulated industries must develop emergency response plans using terms and formats dictated by the regulating agency, which makes these terms and formats proprietary as well. There are many diverse examples of regulated industries that use hazardous or nuclear materials or operate dams.

Many organizations that are involved in emergency response—such as hospitals and emergency medical services, schools, the American Red Cross, and many volunteer agencies, National Guard and military units—may still use other terms and models. Government agencies at the local, state, and federal levels still use a wide variety of security plans and procedures, even though there is now one standard methodology required, as explained in Chapter 23. Most of these procedures are not uniform or comprehensive. And they can be incompatible, present major communications barriers, and cause unnecessary misunderstandings, delays, and wasted resources. In turn, most private organizations are not aware of the government procedures in place or of the many resources that can assist them.

Every organization experiences essentially the same threats and has limited response capabilities and resources. Almost every organization is likely to be overwhelmed in a major emergency, and everyone can benefit enormously from outside resources. Yet, each venue tends to use dissimilar models, terms, and procedures that are often unintelligible to the others. In many cases, their models become file-and-forget plans and procedures that no one understands or accepts.

If only as prudent risk-management strategy to avoid liability issues, a single, uniform, and comprehensive standard for processes, procedures, and language is needed. The best way to do this is to adopt the DHS/FEMA methodology, which encompasses all hazards, coordinates all resources efficiently, and is clear, concise, understood, and accepted by everyone involved. Details are given in Chapter 23. This is now required for all federal government departments and agencies, and for state and local jurisdictions as well, if they wish to continue receiving federal grants and assistance.

A useful public/private standard for evaluating disaster/emergency preparedness is published by the Emergency Management Assessment Program[5] (EMAP), which is a nonprofit professional organization. The EMAP standard covers all the current DHS/FEMA methodology and provides a clear and concise method to determine compliance, as well as a self-auditing process. EMAP also provides an independent, outside security audit of any government organization to ascertain that its emergency management program complies with the federal requirements; it may soon be able to audit private organizations also. The audit is done at the applicant's facility by professional, well-trained volunteers. The whole security-audit process is both rigorous and inexpensive. The EMAP is therefore a suggested alternative to a CPA-provided security audit as well as to the audit procedures of other professional organizations that generally do not include either infrastructure security or compliance with federal procedures.

Security preparedness that is compliant with these standards may become a requisite for the private sector to obtain insurance, to avoid liability and excessive costs, to establish innocence, or to use capital markets for risk management. Knowing these standards will help the threat assessment process and, later, can provide better protection.

22.3 THREAT ASSESSMENT PROCESS. Effective security planning begins with a thorough threat assessment. This process begins by establishing an ad hoc organization, obtaining budget approval and the sponsorship of senior management, and formation of a steering committee that represents all the stakeholders.

The first four tasks are to:

1. Identify all potential threat situations.
2. Determine the likelihood and estimate the direct and indirect costs of each threat.
3. Evaluate and prioritize each threat.
4. Prepare and present a final report that each committee member signs.

This report becomes evidence of due diligence to avoid liability, and becomes the basis for protecting the infrastructure, which is described in Chapter 23.

The way *not* to go through the planning process is for a few people to decide on the threats and risks and then write the security plan. This is the opposite approach to good and effective planning. And it will not work. Very few of the key personnel will accept or even understand the plan, and it will probably be ignored or overlooked when the next emergency arises. Good examples of failed planning are the emergency management plans for New Orleans and Baton Rouge areas that were put in place before Hurricane Katrina in 2005. Apparently, both were good plans. But few officials understood the plans or remembered to use them.

The threat assessment planning process *must* be done thoroughly and completely. For additional perspectives on risk assessment and risk management, see Chapters 62 and 63 of this *Handbook*.

22.3.1 Set Up a Steering Committee. The security planning process is not effective unless all of the stakeholders are represented, and the best way to do this is to establish a steering committee. The committee will help identify and evaluate potential threats, and help develop a comprehensive protection plan.

The committee should represent all stakeholders and include as much experience, knowledge, and perspective as possible. Stakeholders should include users, administrators, management, key partners, customers, vendors, and service providers (such as maintenance, repair, and cleaning personnel). A project manager should participate, as should legal, financial, and human resources representatives. Independent experts and facilitators are recommended as well. The best committee chair is usually an outside facilitator who is immune from political or cultural bias, loyalties, or product preferences. It is also wise to seek input from stockholders, lenders, insurers, and community and government officials.

This can be a virtual committee that rarely needs to meet in full session. Communications by email or phone should be sufficient, and an in-house staff can gather data, issue reports, and work with individual committee members. Confidentiality is not an issue at this point. The committee's purpose is to identify threats and to assist in planning. No member need have full knowledge of the actual vulnerabilities or the resulting security systems and protections. Each member, however, should sign off on the committee's final report.

Once its initial mission is completed, it is best for this committee to convene at least annually to review new and changing threats and to assess how well the security systems have performed. This committee not only represents all concerned stakeholders, but it can provide oversight so that management does not neglect to exercise, review, and update the security plan. In effect, the committee provides due diligence that management has fulfilled its fiduciary responsibilities.

22.3.2 Identify All Possible Threats. The first step is to identify all possible threat situations that might affect the information infrastructure. This is not to say that each threat will someday happen but that, conceivably, it might, nearby or even at a distance. The causal connection could be tenuous at best and the event most unlikely, but it could someday happen. It is far better to think objectively about such things and to plan effectively than to simply write them off as events that could "never happen here" or to claim that, if such a disaster does happen here, "there is nothing we can do to mitigate it." Both statements are patently false and potentially very costly.

The threat list can be very long. The format should tabulate each specific threat with a one-line description to clarify what is meant. For example, "flooding" is too general a term to meaningfully assess its risk. Instead, one might list "riverine flooding," which could be the result of heavy rains, snowmelt, or a breached dam or levee; "local flooding," such as a water-main break, severe storm, or a major fire nearby; or "premises flooding," due perhaps to leaking pipes, drains, or leaking windows, roofs, building setbacks, a burst water tank or cooling tower, or fire suppression. The purpose is to have a long list of specific threat situations that can be individually assessed. There can easily be 100 or more threats listed. As the project proceeds, those involved will surely add still more threats and fine-tune the definitions. The threat list should include threats that are direct as well as indirect and threats affecting the general area, the region, or the whole country. It is best to list all threats, including those that are deemed unlikely to occur, as a precaution against inadvertently omitting some that might have more importance than is apparent at first glance.

Unfortunately, threats tend to cascade: One threat situation can create others, and these, in turn, can create more threat situations. For example, severe flooding can close roadways and keep people away from work, impede delivery of supplies necessary to operate (such as food or water or fuel for an emergency generator), disrupt communications, cause mud slides, ignite fires, trigger looting, or cause civil unrest. As

another example, during a fire there will likely be loss of electrical power for equipment cooling, ventilation, or communications, and released water or chemicals that must be contained. Each of these conditions is a separate threat with its own risk. In reality, few hazards occur in isolation, and many of the cascading events may be unexpected and unpredictable, and can themselves trigger still more events.

Therefore, the tabulation of each threat should also include two other columns: (1) events that could trigger the threat, and (2) cascading threats that could result from the threat. This is easily done by numbering each threat and referencing the numbers that may interconnect the various threats. The information is needed to evaluate the impact of each threat. It can only be a general indication for planning purposes of what might happen. The actual cascading will probably be different, but the analysis of the threats will still be valid.

Do not rely on force majeure—for example, a major earthquake—as an excuse that a particular threat may be unavoidable. Force majeure does not limit liability because such events can be anticipated and some steps taken to minimize injury and damage. Nor may acts of war serve as an excuse, for the same reasons.

Finally, the threat list should be divided into major categories such as natural events, man-made events, vandalism, attacks, support system failures, and so on. The remainder of this chapter builds on the threat list and establishes categories best suited to particular needs.

22.3.3 Sources of Information and Assistance. The next step is a compilation of past events: historical records and details as to what happened, when, and how, what injury and damage resulted, was there warning, how fast was the onset? The compilation should include events in both near and distant areas that could possibly cause damage and disruption, either directly or indirectly. Perhaps the 2005 subway and bus bombings in London or the earlier commuter-train bombings in Madrid could affect American workers who commute and become concerned about their own safety.

As the threat list gets bigger, it is necessary for everyone involved to think out of the box and to look for scenarios that past trends suggest could eventually happen. This thinking can be productive when it examines possible scenarios objectively. The author, as a security consultant in New York, worked with tenants in the World Trade Center and tried to suggest disaster planning but was overruled by the landlord who assured everyone that this was "the safest place on earth."

Information sources to investigate include the weather bureau, libraries, local fire and police departments, state and county officials, utility companies, newspaper files, and local knowledge of past events within the region. Power and communications utility companies can provide outage reports, but their terminology must be clearly understood. For example, an electrical "outage" may only include interruptions lasting more than five minutes. Regional and state regulatory authorities, public utility commissions, industry and professional organizations, and business development groups may also provide useful information, once their perspective and terminology are clarified.

All local, regional, and state emergency management agencies should now have an all-hazard mitigation plan that lists all threats that have occurred throughout their jurisdiction. There should also be dates, descriptions, and map locations in their plan; this information may go back historically for a century or more. And each agency should have a local emergency operations plan that lists all potential threat situations within their jurisdiction. Look particularly at the appendices and annexes that discuss particular types of threats. In addition to what is in their plans, each agency should be of considerable assistance describing possible major threats that are man-made, involve

hazardous materials, or are major health risks. The threats they identify are probably your threats as well.

Also, work with each stakeholder for advice on threats they know about or perceive. Ask each person to reach out to their customers, vendors and suppliers, service providers, business neighbors, and consultants and academics as well. There will be much expertise, perspective, and good advice gleaned from these sources. The time seeking it will be well spent.

At the end of this chapter is a list of some reference books that can be helpful. Although some of the procedures are outdated, there is still much useful source data.

22.3.4 Determine the Likelihood of Each Threat. The steering committee should consider the likelihood that each threat may happen. The best way to estimate likelihood is as an annual probability on a relative scale (for example) from 0 (none) to 5 (very likely). A perceived likelihood of once every 100 years (which equals a 1 percent annual probability) could be rated a 1. A 5 percent annual probability might be rated a 2, while an annual probability of 20 percent or more had best be rated as a 5. Threats from severe weather should be rated higher than historical data suggests in order to account for the world's changing weather patterns. Threats with a likelihood of 0 should also be included, as this value can be used later to show the relative vulnerability of each and every threat.

The steering committee can adjust each rating as it deliberates so that the final results are useful and meaningful.

22.3.5 Approximate the Impact Costs. Once all the possible threats are determined and the likelihood of each is estimated, the next step is for the steering committee to deliberate the impact of each threat. Impact considers the potential losses and costs associated with each threat. (Do not consider cascading threats at this point.) Here also, a relative scale of 1 (very low) to 5 (very high) is suggested. A 0 impact rating is probably not relevant, because, in reality, some response and recovery costs will occur.

There may well be individual factors needed to clarify the importance of the overall impact. Such factors might include the likelihood of injury or death, the amount of property damage, the relative response and recovery costs, and, especially for a private organization, the costs of loss of business. The parameters of each factor and the scope of each rating can be adjusted as the planning proceeds, so that the results are realistic and meaningful. The purpose now is to show which threats are potentially the most dangerous or costly if they occur.

Threat assessment planning is best done using a worst-case scenario because this is what tends to happen in reality. Even then, actual costs incurred can be much higher than anticipated. Although the steering committee has only limited knowledge of possible costs, with some discussion the committee should reach consensus as to which rating to apply to each listed threat. Their findings will be reviewed later by others better able to predict potential costs and in a better position to realize that the impact of some threats is much more costly than others.

The impact costs are both direct and consequential. The direct expenses can include costs to locate the trouble(s), stabilize the systems, repair and install replacement infrastructure, reboot, restore databases, and thoroughly test both the data and the systems. Other direct costs that can result from each event include loss of productivity of system users, overtime needed to regain production schedules, temporary contracted services, and interim facilities, or materials and supplies needed immediately.

Add to these potential indirect or consequential costs such as food, sanitation, and lodging; loss of business, customers, or market share, falling stock price, and the costs of public relations efforts to control rumors and adverse news reports. Costs will continue to accumulate during the response phase, throughout the recovery period, and possibly long after. The totals can far exceed expectations.

22.3.6 Costs of Cascading Events. The possibilities and indeed the likelihood of cascading events have been described. Each event adds to the impact costs, but determining how likely is the cascading, and how much the extra cost, is at best an educated guess. If the causative events for each threat are shown on the threat list and the potential cascading events are there also, the impact costs of each individual event should be estimated. However, cascading may invalidate the likelihood factors. If one threat is already occurring, others may be imminent, regardless of the likelihood shown.

The total impact cost will probably be somewhat less than the sum of the parts. There may be some economies of scale during multiple events, as the response activity for one event can take care of another event for little extra cost. Here is where the experience and perspective of the steering committee, security professionals, and management become valuable. However, the combined costs must still be realistic, or the cost-value analysis described in Chapter 23 will be flawed.

22.3.7 Determine the Vulnerability to Each Threat. Deciding which threats are the most important can be subjective and controversial and based on unproven assumptions. The approach suggested here yields quantifiable data that are realistic and will clearly indicate which threats are the most important and why. Individual opinions will vary considerably. The steering committee reports can yield approximations that are statistically valid—if the rules are followed and everything is done carefully. The committee's deliberations should be reviewed by experts and senior management, using calculations to be described, followed by the cost-value analyses described in Chapter 23.

One other common alternative is a matrix showing likelihood vertically and impact horizontally, rating both factors as high, medium, or low. This yields nine levels of vulnerability for each threat ranging from high-high to low-low. Such an arrangement is neither useful nor realistic. Since allowance need to be made for those impact factors that are potentially more costly than others.

Once everyone settles on the likelihood and impact of each threat, the current vulnerability of each threat can be calculated. Fundamentally, vulnerability is a calculation combining the likelihood of an event with how much damage it might cause. The method suggested yields six levels of vulnerability, from 0 (none) to 5 (very high). The next two steps are oversimplified but can be adapted to most needs. Simply multiplying the likelihood times the impact will yield vulnerability values from 0 to 25, which would not reflect the relative impact costs of most threats. Neither does simple addition, as some models suggest.

1. It is best to use several impact cost factors, as described earlier, and then combine each factor using constant multipliers for each, so that the relative importance of each impact cost is preserved. Combination by averaging each cost factor may not yield useful results. Combining the mean value of each factor (or sometimes the highest value) can yield more realistic data.
2. Finally, convert vulnerability calculations to a 0-to-5-point relative scale by choosing a range of results for each scale value.

Working through this procedure for all threats will show that the formulas, multipliers, and ranges of scale values may need to be adjusted slightly to yield meaningful results. As already mentioned, simple multiplication and/or addition or otherwise linear formulae will not yield realistic results.

The importance of calculating vulnerability is that it can be done in real time, and the multipliers can also be adjusted in real time as threat levels change. For example, this can be reflected in the National Terrorism Advisory System (NTADS)[6] by locale or by possible target types. It can also input warnings from local police and intelligence sources. (Some threat information is probably classified, but the multiplier it can provide is not restricted or even sensitive information.) All the data can be displayed on a spreadsheet, even to the point of color-coding each vulnerability level. Everyone involved in security or emergency management can have real-time access to the same screens, which will update continually as conditions change.

It is well to provide a choice of display screens summarized or in detail by regions and threat categories, or prioritized by vulnerability.

Finally, the complete vulnerability data should be restricted and available only to those with a need to know. This information is highly sensitive; it shows potential wrongdoers where the IS infrastructure is well defended and where it is *not.*

22.3.8 Completing the Threat Assessment Report. Once the detailed threat assessment is completed, a general report should be prepared for circulation to all stakeholders. The report is mostly textual and does not suggest weak spots or what the security defenses may be. Each steering committee member should sign off on this report, and so should the senior management and security staff. The purposes of this report are to develop security awareness, for use in training, and as evidence of due diligence that a thorough threat assessment was indeed performed.

More on how to evaluate and manage the vulnerabilities is described in Chapter 23 of this *Handbook*.

22.4 GENERAL THREATS. A wide array of threat situations can affect the IS infrastructure, degrade productivity, and cause anxiety that will reduce performance and morale. The list suggested earlier begins to identify some of these threats. Some other possible situations that are not generally associated with the IS infrastructure are listed in this section. However, the various threats suggested in this chapter are far from a complete list. Therefore, expert advice, local knowledge, and analysis of current events should all be utilized to customize the threat list to the particular needs of each jurisdiction and then to review and revise the list periodically.

As mentioned before, specific threats are usually divided into broad categories that can easily be summarized to provide situational awareness. Generally, threats are categorized as natural hazards, technical and man-made threats, civil unrest, vandalism, and attacks. Other less likely but potentially more serious possibilities include bombs, hazardous materials release, a pandemic, and threats relating to weapons of mass destruction. A final grouping could be local events that are specific to the sites being protected. The groupings themselves are not important as long as all possible scenarios are covered. Some allocations will be arbitrary, as some threats could be put into any of several categories. (But never list any specific threat more than once.) There is as yet no single, uniform, comprehensive threat list format. As the planning process continues, the best threat groupings will likely evolve.

22.4.1 Natural Hazards. Natural hazard events are becoming ever more frequent, damaging, and widespread, as their patterns seem to be changing. Increasingly recent events include major flooding, severe thunderstorms; hurricanes and tornadoes; blizzards, heavy snowfall, and ice storms; wildfires and heavy smoke; contamination of air, water, buildings, or soil; and the increasing threats of disease. Although earthquakes tend to occur in cycles, these too seem more prevalent currently. Any of these can disrupt business in any number of direct and consequential ways—many unanticipated—depending on the locations of information systems, networks, terminals, data storage, cables, and utilities. Natural hazard events cannot be prevented, but their impact can be mitigated. A closer look at each category follows.

- **Atmospheric hazards**. These can include severe weather events, such as tropical cyclones and hurricanes; severe thunderstorms with hurricane winds, strong lightning, and large hailstones; tornadoes; windstorms; blizzards and heavy snows; ice storms; air pollution and high ozone levels; nuclear fallout (which is always occurring, but at low levels); extreme cold weather; and extreme hot weather.
- **Geologic hazards**. These are mainly landslides and mudslides, land subsidence (sinkholes), and expansive soils due to water.
- **Hydrologic hazards**. These hazards include riverine flooding and flooding of low-lying areas due to heavy or prolonged rains; rapid ice or snowmelt; ice jams or debris obstructions in waterways; coastal damage due to storm surge; erosion of streams; and collapse of roadways, bridges, and buildings. Hydrologic events can cascade to disrupt water and sewage systems, cause food and fuel shortages, and even trigger fires or explosions.

  A dam or levee failure is a major hazard that must be carefully considered and then reviewed with experts. Most dams are well maintained and protected from natural hazards but not from vandalism or terrorist attack. Nor are most reservoirs well protected from chemical attack. The consequences of such events can be far more devastating than imagined. The 2005 failure of levees in New Orleans caused by Hurricane Katrina produced catastrophic results over wide areas.

  A prolonged drought can be particularly disruptive. There may be many fires and heavy smoke that require evacuation. Potable and cooling water may be in short supply. Both people and equipment must be protected from dust or smoke that cannot be done effectively within many facilities. Hydroelectric power may be rationed and rolling blackouts imposed.
- **Seismic hazards**. Such hazards include earthquakes and tsunamis. Many areas of the United States are at moderate or even high risk of a large earthquake. There were major events in these areas hundreds of years ago, and it is believed the time is nearing for a recurrence. For example, the Ossipee Fault in central New Hampshire produced a "San Francisco–size" earthquake about 300 years ago and may occur yet again, which would cause significant damage to Boston and its suburbs. At the least it would topple equipment cabinets, break cables and wires, and probably disrupt cooling and ventilation to IS systems.

  Tsunamis, which are caused by undersea earthquakes, could at some time hit U.S. shores, based on historic events. And like Hurricane Katrina, a major seismic event anywhere in North America could impact businesses throughout a wide area.

22.4.2 Other Natural Hazards

- **Major volcanic eruptions**. Volcanic eruptions can spread atmospheric dust worldwide, which can affect weather and cause severe storms, damage cooling systems, disrupt radio and satellite transmissions, and restrict air travel.
- **Wildfire**. Fire can close transportations routes, disrupt utilities and support systems over a wide area, require evacuations, and create heavy smoke that is dangerous to people and equipment.
- **Blight or infestation**. Caused by disease, weather, or insects, blight or infestation can create health problems and disrupt food supplies, which can indirectly impact business.
- **Sunspot activity**. The approximately 11-year cycle of solar magnetic storms that cause sunspots may cause waves of electromagnetic radiation that can disrupt communications and even the electrical power grid.

22.4.3 Health Threats. Of increasing concern is the possibility of health emergencies, such as a SARS or an anthrax outbreak, or West Nile virus in relatively small areas. Such outbreaks can probably be contained with antiviral medicines and vaccines. However, business disruptions will surely result. Of greater concern is the threat of another pandemic that could quickly spread through large populated areas. As yet, there are no effective means to prevent either the onset or the spread of a pandemic, other than isolation and quarantine, which includes closing businesses and schools and perhaps many municipal offices as well. As mentioned earlier as a possible worst-case scenario, personnel refusing or unable to report to work during a major health emergency might total as much as 40 percent of the workforce, and they might remain away from the workplace for as long as 14 months. Consult local health officials how to access detailed information about these threats.

22.4.4 Man-Made Threats. *Accidents* cause most security problems. Accidents, sloppiness, and blunders are the consequences of bad design, poor quality control, improper installation, failure to update, and poor maintenance. Disruption often occurs during maintenance. Deliberate actions such as snooping, pranks, vandalism, and spying are all increasingly prevalent and sophisticated but are still overshadowed by accidental, unintended events.

Moving furniture or equipment (unless done by trained professionals) can damage wiring, connectors, and other equipment sufficiently to crash systems. And while it rarely happens, substantial "accidental" damage can occur during a labor dispute or when nonunion personnel are brought on site.

Construction work, alterations, repairs, and wiring changes often cause damage to information systems. Crews often drape drop cloths over workstations and equipment to keep off dust and debris. But no one thinks to shut down the equipment first, so it becomes overheated and probably fails either immediately or soon after. Crews also plug in power tools, floor waxing machines, or vacuum cleaners in whatever electrical outlets are handy. If these happen to be dedicated circuits for systems equipment, damage may well result. In like manner, workstation users often plug time stamps, electric staplers, refrigerators, fans, and immersion heaters into outlets intended only for information systems. Such mistakes can cause intermittent problems that are difficult to locate.

Wiring runs and exposed wire can also create threats. Wiring and cabling are vulnerable in many ways—and are becoming increasingly more so. Data cables tend to be fragile and easily damaged by accident, moving furniture, cleaning and maintenance, improper connections (such as unauthorized equipment), rage, or deliberate attack. Metal wires are also vulnerable to electrical and magnetic interference from nearby light fixtures, motors, transformers, or RF emitters, as will be explained. Lightning is attracted to, and electrical power surges may be induced onto, any metal wires.

Fiber-optic wiring, which is fast replacing metal connections high-speed communications, is especially fragile. Any unusual pressure or a sharp bend in any fiber cable can alter its transmission characteristics and may cause failure. However, fiber is not affected by any electrical or RF inference.

Cables, patch panels, cords, and connectors are often exposed to accidental damage during routine premises maintenance and cleaning. Vacuum cleaners and shampooing and waxing machines can easily damage both signal and power wiring. Carpet shampooing can soak connections and flood underfloor wiring. The casters on chairs and equipment or the moving of furniture often damages cables and connectors.

Intrusions are a major threat. There are many ways to gain access to the infrastructure of information systems. Once an intrusion is successful, repeated disruptions may follow due to accidents, mistakes, snooping, hacking, spying, vandalism, extortion, or deliberate attack. Intrusion is not usually prevented by premises access control. The infrastructure's vulnerability points are different, and early detection is necessary to prevent trouble before it happens.

Intrusions via wiring and cabling are covered in the prior section. Here the emphasis is on preventing access to hardware, distribution and termination panels, patch panels, or any of the utilities that support them. Assessing the possible threats first involves inspecting the existing infrastructure and the defenses already in place. It is also necessary to determine what physical threats and vulnerability points are not sufficiently covered by premises and logical security.

Unauthorized persons should not be allowed access to any type of equipment room, rack, or closet. Everyone granted access, including visitors, should be logged in and out. And as a second layer of protection for critical areas, a gatekeeper should be at hand to observe everyone who enters and leaves, to authenticate each person's identity, and to know why access is needed. Gatekeepers may include guards, receptionists, supervisors, or managers. Except where high security is required, the gatekeeping functions can be performed remotely using surveillance and access control systems, and monitored by motion and proximity detectors.

Workstations located outside of equipment rooms are generally protected logically—using login procedures, tokens such as smart cards, and biometric devices—depending on the sensitivity of the data the system handles. Theft of equipment, components, or removable media is possible, especially during nonbusiness hours. Substitution, rather than theft, is a much more serious threat. The only safe defense is that all data must always be fully encrypted. Should a theft occur, spare units should be immediately available from nearby secure storage. To save time, thieves often cut signal cables rather than disconnect them. It is therefore wise to keep spare cables in a nearby secure storage.

There are various theft-deterrent devices to clamp or tether equipment to nearby partitions or furniture. These devices can be cumbersome and intrusive, may hinder maintenance or repairs, and most are quickly defeated with bolt cutters or a pry bar. Intrusion alarms are more effective but only detect an attempted intrusion rather than deter theft. Intrusion alarms are usually silent and triggered whenever a case or cabinet

is opened. Software can trigger alarms also when a cable is disconnected or broken, or when equipment is removed or fails. Good, sturdy equipment and cabinet locks are more effective and can dissuade theft, provided the equipment and cabinets themselves are not easily removed.

22.4.5 Wiretaps. Wiretaps are another means of intrusion whose purpose is to copy data and sometimes also to change it surreptitiously. Most taps can be placed quickly and inconspicuously even in occupied office spaces. And most wiretaps are very hard to detect—if not impossible—except by close visual inspection of all possible entry points, which is time consuming and cumbersome, and must be repeated periodically. Very often, wiretaps are simply not noticed.

Taps are used for surveillance and spying, changing or erasing data, stealing software, theft of service, injecting worms or viruses, or planting decoys to trick responders into thinking they have found the true cause of an incident. Any physical access to the interconnect points, cable runs, servers, clients, or network equipment is a major vulnerability. An experienced person can place a tap within a couple of minutes, unobtrusively, even when under escort and closely watched.

A wiretap is basically a monitoring device. It may be a small box connected directly into a circuit to tap off and record its data. The data can then be retrieved manually or be transmitted by wire or radio signal to a remote location where the information can be permanently stored and analyzed. The monitoring device may be a splitter that can tap off some of the signals traveling through the circuit and reroute them to a removed location or to a small transmitter that is close by and well concealed. The better tapping equipment uses system power, not batteries, so its service life is unlimited. Monitoring the data obtained can occur within the building or outside, via the Internet, from a nearby vehicle, or via a telephone connection to anywhere in the world.

Wiretaps have long been illegal in the United States without a court order. Years ago, an illicit phone tap in New York City, for example, was patched to a leased telephone line to Mexico where the actual surveillance could take place legally. The patch connection was made by someone at a telephone company central office and was not likely to be discovered. But if it eventually was found, the network could be traced and disconnected. This procedure suggests the lengths and expense that a determined adversary will go to, and it also suggests the high value of the information that could be gleaned. This was years ago, long before today's easy access to the Internet, cheap long-distance calls, Wi-Fi, microwave, satellites, and other broadband systems.

Fiber optic circuits are difficult and expensive to tap and sometimes impossible without breaking the circuit or changing their parameters, which can be detected quickly. There is new high-tech gear available for optical taps. But even with this, good results are problematic.

The first method is to create a sharp bend in the cable and detect some of the light rays that may leak from the outside of the bend. The signal is at best very weak. Cable can be constructed to prevent this. Or cables can be run through metal conduit to make access and then creating a sharp bend very difficult, even at junction boxes. Last, a sharp bend will change the optical impedance of the cable run and will slightly degrade its performance, perhaps enough to effectively crash the circuit. The change in optical impedance can be measured, including the approximate location of the anomaly. The cable-bending method may work, but it is hard to accomplish, detectable, and readily observed by inspection. This method can monitor data, but it cannot inject data.

The only direct method to tap a fiber cable is to open a connection, quickly interpose an interface device, and then reconnect the circuit. Alarms should be triggered

immediately when a circuit breaks or otherwise fails, and these may also give some idea where the problem occurred. But with a good interface device and by working quickly, the tap can be inserted and conditions restored to look normal before anyone can respond. If this method succeeds, the intercept quality will be excellent, and data can be injected into the tapped line as well.

Spying on what is carried by fiber circuits may be possible from within equipment rooms or at other points where the optical data are converted to electronic form or to radio waves. The equipment to do this and its capabilities are highly classified and generally unknown to most security experts.

Wireless transmissions do not need to be tapped, but the process still applies. All that is needed is a very sensitive radio and a high-gain antenna, often situated in an outside vehicle. Chapter 33 in this *Handbook* discusses wireless-network security in detail.

Metal wires are relatively easy to tap into undetectably. Small, inconspicuous induction coils placed next to a wire can easily pick up the signals without penetrating the insulation. Sometimes the taps are done with tiny probes that reach around the circumference of a wire. Induction taps are generally undetectable, except visually. However, induction taps are difficult when multiconductor, twisted-pair cables are used, especially those where each pair is shielded and the entire cable is metal clad, as are some aerial telephone cables to protect them from radio-frequency interference.

When induction is not practical, bridged taps easily placed within an accessible junction box, cross-connect or patch panel, test box, or termination strip. Bridged taps can be employ a physical connection. On long cable runs, usually one or more inline terminal strips interconnect two cables. These are good locations for taps, especially where there may be unused wire pairs within the cable itself that can be used to route the tapped data to a safer location. Bridged wiretaps can also be inserted anywhere along a wire by penetrating the insulation with a needle or cutting away a tiny piece of the cladding and insulation in order to splice in a tap wire. Bridged taps are usually very high impedance, so they do not alter the circuit; therefore, like induction taps, they are undetectable, except visually. However, skill is required to place them to avoid damaging the tapped wires. As with fiber wiretaps, there are also high-tech methods of monitoring what metal wires carry, the details of which are classified.

Today, the biggest threat from wiretaps is that they are usually done by foreign intelligence, organized crime, or terrorist groups that are very well equipped, highly trained and experienced, and have almost unlimited funding to carry out their work. It is high likely that most of their surveillance is not detected, located, or removed until long afterward—if ever.

22.4.6 High-Energy Radio-Frequency Threats. A controversial area of research involves high-energy radio-frequency (HERF) weapons, which some experts claim pose a serious threat to disrupt, damage, or destroy electrical and electronic systems and equipment over a very wide area. There are several scenarios.

A prototype HERF weapon demonstration at security conferences some years ago showed that a bulky apparatus constructed from easily obtained components can cause personal computers to malfunction while the device continues to operate. It can perhaps cause some system failures as well. There has been no more public information since then as to whether a smaller, portable version could emit a radio-frequency beam of sufficient power to disrupt systems dozens of meters away. Such a device is possible but may not be very practicable, except possibly from a large vehicle.

Along with this is another threat called high-intensity radio-frequency flux, which is simply interference from other electronic devices. These can be nearby transmitters, such as high-powered radio, TV, or radar systems, or transmitters in vehicles. Or there may be interference from close-by RF emitters including cell phones, computers, and devices designed for this purpose. Some of these possibilities are design issues, but when portable RF emitters are involved, the situations become infrastructure threats.

In 2004, another approach was reported that is potentially able to destroy unprotected electric and electronic system nationwide by way of an electromagnetic pulse attack (EMP). This report[7] comes from a commission finding that a low-level nuclear explosion by a missile in the upper atmosphere could cause a massive electromagnetic pulse wave that could wreak catastrophic damage throughout most of the United States. Without extensive shielding, electrical and electronic systems could be destroyed, and power transmission systems as well. However, it appears that such an explosion must be done just right or the damage would be minimal.

The defense against all such RF threats is to shield equipment and wiring effectively. A Faraday cage can envelop and protect everything. This is usually one or more layers of copper screening that are well grounded. There also are Faraday bags to protect and store small items like cell phones or circuit boards.

However, once shielding is installed, it is wise to test that it actually works as expected. The author once toured a large, new data center of a major money-center bank. In touting the huge sums spent on extensive security protections for the new facility, someone mentioned that an impenetrable Faraday shield enveloped the entire room. The only problem was that a transistor radio on top of a cabinet was playing clearly and loudly. Had the room actually been well shielded, no broadcast signal could have penetrated.

Hazardous material incidents can include a release within a building or industrial plant, or from an aircraft, railroad, highway, or waterway accident. Many hazardous substances are routinely transported, stored, and processed throughout the United States. Many of these materials are extremely dangerous when released. They can quickly affect a wide area and may require immediate evacuation. In addition to accidental or deliberate release, these materials may be stolen to become part of an attack elsewhere. And those who attempt to steal, transport, and process the materials into a weapon may endanger others accidentally along the way.

Toxic threats include a buildup of radon within occupied space or in a water supply. Radon is released naturally by many types of rock. It is very prevalent and extremely dangerous if it can enter the body. Radon is an alpha particle that cannot be detected by most traditional measurement devices. Similarly, many toxic substances can accumulate in drinking water or building air. Any of these substances will be disruptive and cause anxiety far beyond the affected areas. Consult local health officials to identify and assess these threats.

Fire and smoke from remote events can be troublesome in several ways. Smoke from events hundreds of miles distant can cause health problems, delay supplies and deliveries, and cause food or fuel shortages. Workers may be unable to get to or from work, and may also be concerned about their families, friends, and homes. Electronic equipment is particularly susceptible to environmental damage from smoke or other airborne particles. And there may be utility problems or loss of cooling caused by a large, distant fire.

Nearby fire or smoke can be particularly troublesome. In addition to the problems just mentioned, an immediate system shutdown and personnel evacuation may be required. Power and data lines may be broken or damaged by heat or water. There may

be flooding as well, and hazardous materials released that are potentially injurious to people and equipment. There will also be the need to contain the threats and to clean up the environment before full business can be resumed.

Smoke, dust, or other airborne particles can cause equipment failures, as these block cooling systems, clog filters, and build up inside equipment, constricting ventilation and convection cooling. Although these are mostly maintenance issues, they are also potential threats.

22.5 WORKPLACE VIOLENCE AND TERRORISM. The use of force, harassment, or physical violence is an increasing reality within the workplace. Drug- and alcohol-related incidents are on the upswing, as are workplace crimes necessary to support and conceal such habits. Adding to the threats is the increasing presence of rage within or upon the workplace. Any of these incidents can cause widespread trauma and disrupt business operations for months. Actual violence or terrorism, or simply the threat or the fear of this, can be extremely costly and disruptive.

Any violence situation—whether threatened, imagined, actual, or peripheral—can seriously disrupt information systems, whether the infrastructure is actually in danger or not. Performance, productivity, and morale will all plummet and remain low for a long time following any perceived or actual threat. Full recovery can take many months—assuming that more incidents do not occur in the meantime. Therefore, a safe working environment, good security planning and implementation, and security training exercises are essential so that everyone feels safe.

The likely tools of choice for inflicting widespread injury and damage may soon be weapons of mass destruction (WMDs) rather than the old-fashioned knives, guns, bombs, or arson that inflict only limited damage. WMD devices are far more dangerous, and many are small and easily concealed in a pocket, package, or briefcase.

A small, common-looking object containing biological toxins and powered by a flashlight battery can theoretically kill every person within the largest of office buildings. A vial no larger than a lipstick can contain enough virulent hemolytic viruses to kill every person within a 20- to 50-mile radius, if it is dispersed efficiently. Because few chemical or biological WMD compounds have much odor or color when they are disbursed, occupants, visitors, bystanders, and responders are all likely to be innocent victims. Illness can begin within minutes, hours, or days. Laboratory analysis is needed to identify many of the substances. More time is needed to determine the scope and spread of the damage. Ordinary personal protective gear and breathing apparatus can provide little or no protection.

WMDs can be enormously destructive. A national FEMA, Department of Justice training exercise called TopOff that I conducted in May 2000, simulating a biological attack on Denver, Colorado, resulted in an estimate of 57,000 fatalities. Since then, similar exercises have not released fatality estimates.

WMD threats can be grouped by the mnemonic B-NICE. They include *biological* agents, such as anthrax, cholera, pneumonic plague, tularemia, Q fever, Ebola, smallpox, botulism, ricin, and some others. These are all living bacteria whose incubation periods (the time of onset following exposure) are measured in hours or days.

Nuclear and radiological releases can cause widespread panic but are not likely to cause mass casualties, except to persons nearby. Radiation levels can easily be monitored, and building structures will often serve as an effective shelter.

Incendiary devices are utilized mainly to cause structural fires. A device can be planted surreptitiously and then triggered remotely or by a timer. Rockets and small

missiles are increasingly an incendiary threat, as are 9/11-type incidents: airplanes, vehicles, or boats used as bombs.

Explosive devices are similar. These may be stolen or smuggled ordinance or, increasingly, improvised explosive devices made from commonly available materials.

Combined attacks are possible, such as explosive devices used to disburse chemical agents (but not biologicals, which are live bacteria that would be killed by an explosion) or the so-called dirty bombs, which are explosives used to disburse radioactive material.

In a worst-case scenario, all WMD events can be exceedingly dangerous and disruptive, and any of these events may someday happen within the United States. Even the rumor of a WMD event can be damaging and cause widespread hysteria, panic, and create a large army of the walking well. Whether the threat is real or imagined, WMDs are costly to combat.

Most details on WMD and extremist or terrorist threats are classified, but state and federal authorities should be able to provide some insight for an effective threat assessment. The authorities should at least be asked about perceived threats, possible target areas, and regional and local incidents and security concerns.

22.6 OTHER THREAT SITUATIONS

22.6.1 Leaks, Temperature, and Humidity. Threats involving water and other liquids that may be hazardous should be considered also, as well as temperature and humidly conditions where equipment is located. Sprinkler systems in nearby spaces can cause equipment damage, as can liquid leaks from storage tanks, cooling towers, or pipes near equipment areas. Atmospheric conditions near equipment are threat situations as well. Air temperatures that are either too high or too low can cause equipment failure. High humidity can cause condensation within equipment and, worse, a form of galvanic action that degrades connectors that will then fail eventually. Also, low humidity is a threat because this promotes static electrical discharges that can be deadly to electronic gear and is often undetected until the equipment fails without warning. (See also Chapter 23, Section 23.8.7.)

22.6.2 Off-Hour Visitors. Cleaning and maintenance personnel usually work off hours and are often hired by a contractor or landlord who rarely provides much, if any, background checking, supervision, or training. Very few of the personnel are aware of security precautions and most know very little about the IS systems their work can damage. Many are paid poorly, are forced to rush their work, and may understand little English.

One of the classic disruptions can occur when workstations and associated devices are simply plugged into wall sockets or extension cords without labels indicating that they must not be unplugged; cleaning staff may easily unplug such devices in all innocence to run their own cleaning equipment and never even notice that they have abruptly powered down the computing equipment on nearby desks.

Waxing floors and shampooing carpets are usually done off hours by outside services. Moving furniture and changing workstations are often done after hours, as are repairs, alterations in occupied office space, and other major maintenance. Almost always, these people are unescorted, and many are not even logged into or out of the premises or identified in any way. Worse yet, many of these people prop open doors so they can work faster and sometimes so they can take advantage of air-conditioning in adjacent spaces.

Persons who are unknown often come into the premises without proper authorization. Some work for the landlord or the service organization. Some are delivering food. Some are messengers. And some are snooping, spying, looking to steal, or perhaps bent on violence. Even daytime workers may show up after hours. Therefore, the best security policy is to admit no one after hours until identified positively, logged in, and a need to be there is established. Everyone should be stopped at a reception or delivery desk, with no further access into the workplace until properly cleared. Anyone leaving should be logged out as well, especially if they have been out of sight of where they first entered. Screening visitors at any hour has the added security benefit that outsiders do not see where any IS infrastructure may be located or where money, wallets, or handbags are kept. A night bell or intercom outside an always-locked door should be used to prevent entry to a sensitive facility.

22.6.3 Cleaning and Maintenance Threats. A floor-waxing machine abrades everything it touches and will very quickly destroy unprotected wiring, connectors, dangling cords, or unseen power extension cords (which, by the way, are illegal according to most electrical codes). The waxing machine operator often cannot see any of these items or may not have time to look carefully at what is plainly visible. Carpet shampooing uses a lot of liquid and can flood floor-level outlet boxes and drain into underfloor ducts and conduits. Electrical plugs, receptacles, and unauthorized extension cords for critical equipment are often unlabeled. The cleaning staff does not know they are critical and can unknowingly unplug servers to power their cleaning equipment. Therefore, the threat assessment must first determine whether premises design invites problems, even though cleaning and maintenance personnel do their work carefully.

Users often compound these threats. Many workstations may be logged off but are left running continuously. No one tells users to shut down and cover their equipment before major cleaning or maintenance.

22.6.4 Storage-Room Threats. Rooms used to store computer supplies, paper, or forms are especially dangerous if a fire occurs. For example, stored cartons of paper, forms, or stationery expand when they burn, then burst into a conflagration that spreads quickly and burns at a very high temperature. Such a fire occurred in a high-rise office building in Manhattan. Even though the fire department quickly contained the fire, some steel building columns were so weakened by the heat that the building nearly collapsed. The cause of this fire was determined to be a cigarette butt that fell between stacked cartons of paper. This was assumed to have been accidental but could well have been arson.

Any storage room containing sensitive or inflammable materials must have smoke detectors, sprinklers, or other approved fire suppression systems designed to protect both the room and its contents. There must also be fire extinguishers nearby. Storage rooms should be kept locked, and access should be limited to trusted persons, if only to protect the value of the contents. There should be access control systems with admittance limited to authorized persons only, with open-door alarms provided as well. Delivery persons should always be continuously escorted. Loitering, smoking, drinking, drugs, snoozing, or any social activities must be kept out of and away from all storage areas.

22.6.5 Medical Emergencies. Most medical emergencies will cause business disruptions. People stop work to see what is happening, and if they know the victim

they may remain demoralized and unproductive for weeks afterward. Many will fixate on the circumstances and increasingly involve others to compound the problems. Even minor medical emergencies can cause large and long disruptions. In many medical emergencies, the first five minutes can decide life or death, and professional medical assistance is very likely unavailable that quickly. Any death that occurs in the workplace will result in long-term, widespread, and lasting trauma and disruption. (See also Chapter 23, Section 23.9.4.)

There must be first aid supplies, oxygen, and automatic electric defibrillators (AEDs) handy in every workplace. And there must be people nearby who are trained and currently certified in cardiovascular resuscitation and first aid. The costs of a disruption alone are far greater than those for providing ample medical supplies, equipment, and training. In addition, a first aid room and a trained nurse are wise precautions and probably a cost savings as well.

Prompt medical attention is essential for senior executives, visitors, and all workers and visitors. Medical threats will happen and will be very costly if medical assistance is not immediately available.

22.6.6 Illicit Workstation. A convenient method to set up a logical intrusion or attack is to unplug a desktop terminal or workstation that has limited functionality and substitute a full-featured machine well programmed with spyware and analytic utilities. Anyone with a full-featured notebook computer and physical access to the network may be able to connect easily.

Illicit users may then be able to log onto the network, possibly entering the user's own trusted password. They could then search for restricted information and use the full-featured machine to copy network data. Then the illicit user simply connects to a nearby telephone receptacle to export network data via a dial-up connection. (An Internet connection would likely be blocked by the firewall.) The illicit user may use hacking programs to gain supervisory status and spyware to access more sensitive information, crack passwords, steal software, or infect the network with malicious software. The illicit user may install a backdoor entrance into the network that an accomplice can use to spy, monitor network traffic, and modify or destroy data. An experienced user may then erase all evidence of any intrusion.

The intrusion might be done by an in-house employee, contractor, service technician, vendor, or consultant who could breach the network while ostensibly checking a user's machine or LAN connection. It could also be done within a hot zone using a Wi-Fi connection where there may be few, if any, firewall or security protections. Maintenance personnel working off hours also could substitute a terminal inconspicuously. However it is done, intrusion is a serious physical threat.

Security awareness is the best prevention. No one should be swapping equipment or connections unless a manager, supervisor, or nearby workers know the person's identity and what he or she is doing. If there is any possible doubt, the activity should be reported. The second best prevention is good network security, with alarms when any desktop systems is opened, disconnected, or shut down.

Chapter 15 in this *Handbook* includes detailed discussion of penetration techniques.

22.6.7 Other Local Threats. There are many threat situations specific to a group, organization, or community. Usually, many of these situations are identified and assessed as planning progresses. For example, for a community, important local

threats might include vandalism or actual damage to school buildings, schoolyards, or school buses, a building or bridge collapse, interrupted power or energy transmissions systems, damaged fuel storage facilities, or a communications failure. Many possible specific situations will likely be discussed during the planning process, while review and advice by outside experts and government officials probably will lead to more important situations. Some considerations for local threat situations include:

- **Utility disruptions**. Disruptions happen more frequently than reported. Power outages can last a few minutes, hours, or days, caused by storm damage, equipment failure, tree branches hitting power lines, highway accidents that topple utility poles, digging accidents that break transmissions lines, and vandalism or worse, such as toppling transmission line towers. There are increasing reports of transmission towers and utility poles deliberately toppled, power outages caused by insulators smashed or wires cut by bullets, and outages that occurred when wires and bus bars were stolen for the value of the metal in them. There can be static or RF interference or spikes on a utility line at any time due to atmospheric conditions, lightning, or the utility switching their feeders.

  Communications outages can be even more problematic, and few of these situations are widely known. Communications providers are for-profit enterprises and cannot provide much fail-safe protection, backup, or redundancy when severe weather, vandalism, deliberate attacks, or equipment failure interrupt their services. Although there not many alternate suppliers to choose from, the only security protections are redundancy and alternate suppliers.
- **Civil, political, and economic disruptions**. Such disruptions may be indicated by an elevated Homeland Security Threat Level or by state or regional alerts. Other possible disruptive events include a demonstration, march, disorderly group, or an unruly crowd. Other economic emergencies include a plant closing, a strike or a lockout, a transportation failure, or a shutdown. There are also the threats of violence of any kind; a hostage incident or kidnapping; sabotage of any infrastructure (e.g., power lines); contamination of food, water, air, or soil; a food or fuel shortage, a spike in energy costs or a shortage; and the repercussions of a major evacuation somewhere in the region.
- **Coordinated attacks**. These attacks are also possible, perhaps even by terrorists. Here, many points of the infrastructure are attacked at once, and many forms of attack may be used. There may be diversions in order to plant surveillance devices, place stronger weapons, or simply to distract response teams and stretch their resources more thinly. And whether spying is intended or not, the goal of a coordinated attack is to inflict maximum damage and disruption.
- **Solar storms**. Solar activity can also cause large problems. Sunspots affect electrical distribution systems as well as electronic systems, and they can severely disrupt satellite, microwave, and emergency communications, and radio, TV, and radar transmissions.

22.7 CONFIDENTIAL THREAT INFORMATION. Many other threat scenarios should not be described publicly because they are easily done by anyone with a grievance, and they tend to incite copycats. Other threats that will not be described utilize simple tools or devices that are readily obtained, are inconspicuous while in use, and can be safely hidden after a crime. Avoiding such threats is difficult, and

apprehension is unlikely. Nonetheless, mitigation is possible, and sometimes even deterrence, once these threats are known and understood.

Most vendors, installers, and consultants have long lists of such unmentionable threats and ways they know of to disrupt, snoop into, or destroy specific types of information systems. No threat assessment can be complete without asking all these information sources for whatever attack methods they can suggest as well as for methods of preventing or mitigating each threat.

Many useful sources of information and guidance are not generally public—many sources, indeed, that Internet search engines have not discovered. Some of this material is simply not publicly available, and other sources may be derived from classified documents. But briefings or redacted information may be made available to those who need it. More information may be available from local, state, and federal authorities, regulatory bodies, peer groups, business or vendor associations, professional groups, and security experts and consultants. Each of these sources may be willing to share information not available to the general public.[8]

The FBI is the ultimate source of threat information. But this information is mostly classified, and without the proper security clearance and a need to know, the FBI will not divulge much. A nationwide group of people involved in IS security is sponsored by the FBI and does provide some useful (sensitive but unclassified) information, once each member is vetted and approved. This group is called InfraGard, and there are chapters in most states. Visit www.infragard.net for a list of local chapters and contact them about membership.

22.8 SUMMARY. A vast array of possible physical threat situations can disrupt the infrastructure of information systems and thereby also disrupt business productivity and performance. The list of specific threats may identify several hundred situations, and each one should be included and considered during a threat assessment process that is well done and thorough.

Some of the threat situations are obvious and well known, while many others are much less so. Many are newly emerging, suggested by recent events worldwide and by reexamination of historic data. There also are issues of apathy, denial, and ignorance, where some persons feel that such things will never affect them or that nothing can be done to prevent a disaster. Both assertions are patently false and potentially very costly if pursued. Most threat situations can and will happen somewhere, someday. But whether they do or not is a matter of statistics and not conjecture.

Every possible threat situation can be mitigated to some extent by careful security planning and concerned management to minimize injury and damage. In fact, good physical security can be affordable, effective, and efficient. The first step is a thorough threat assessment to identify and consider all possible threat situations. The next step is to determine the likelihood that each threat may occur, the potential impact if it does, and the vulnerability of the organization within the context of its current security protections. Each step can be calculated statistically to determine the best possible options.

This chapter begins a threat assessment process that takes all these factors into account. Chapter 23 then suggests some ways to protect the information infrastructure and completes the threat assessment process with a cost-benefit approximation to determine the vulnerability of each threat.

These chapters are based on a uniform and comprehensive methodology recommended by the federal government and now required for all federal, state, and local

agencies and departments. Compliance is strongly recommended for the private sector, if only as an effective means of risk management.

The value of this approach is that good security is a wise investment. Anything less than good security is a waste of time and money.

22.9 FURTHER READING

Alibek, K., T. Dashiell, A. Dwyer, and S. Lane. *Chem-Bio Handbook*, 3rd ed. Jane's Information Group (ASIN B008DM4KT4), 2005.

Baker, P. R., and D. J. Benny. *The Complete Guide to Physical Security*. Auerbach (ISBN 978-1420099638), 2012, 360 pp.

BSI. *IT-Grundschutz Catalogs*. German Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik). English version, 2012. https://www.bsi.bund.de/EN/Topics/ITGrundschutz/ITGrundschutzHome/itgrundschutzhome_node.html

DoT. *Emergency Response Guidebook* (2012). U. S. Department of Transportation, Pipeline and Hazardous Materials Safety Administration. A useful quick reference to the characteristic of hazardous materials, including some WMDs, with a table of isolation and initial-response guides. Free copies may be downloaded from www.phmsa.dot.gov/hazmat/library/erg

Fennelly, L. *Effective Physical Security*, 4th ed. Butterworth-Heinemann (ISBN 978-0124158924), 2012, 384 pp. FEMA. *Multi-Hazard Identification & Risk Assessment (MHIRA).* U.S. Federal Emergency Management Agency. Publication 9-0350 released in 1997. This in an excellent reference manual, 355 pages. All types of natural hazards are described in detail, and maps show the likelihood of each threat throughout the country. See www.fema.gov/library/viewRecord.do?id=2214

FEMA. *State and Local Mitigation Planning, How-to Guide*. U.S. Federal Emergency Management Agency. There are nine manuals, FEMA Publications 386-1 through 386-9. www.fema.gov/hazard-mitigation-planning-resources#1

22.10 NOTES

1. See, for example, Duncan Chappell and Vittorio di Martino, *Violence at Work*, 3rd ed. (Geneva: International Labour Organization, 2006). Review available from www.workplaceviolence911.com/docs/20060624.htm
2. The International Labour Organization of the United Nations provides a good view of the rise of workplace violence worldwide. For example, look at the paper at www.ilo.org/public/english/protection/safework/violence/violwk/violwk.htm
3. www.cesnur.org/testi/FBI_004.htm. Every U.S. police department received a copy of the Project Megiddo report and may provide access.
4. The National Institute of Standards and Technology (NIST) within the Department of Commerce publishes many standards and procedures that relate to computer security. The current list, as of November, 2013, at csrc.nist.gov/publications/index.html
5. EMAP, P.O. Box 11910, Lexington, KY 40578, Tel: 859-244-8222; www.emaponline.org
6. See www.dhs.gov/national-terrorism-advisory-system

7. Report of the Commission to Assess the Threat to the United States from Electromagnetic Pulse Attack. Available at www.globalsecurity.org/wmd/library/congress/2004_r/04-07-22emp.pdf
8. Multi-State Information Sharing and Analysis Centers (www.msisac.org/about). Regional Information Sharing System Automated Trusted Information Exchange (www.riss.net). DHS Daily Open Source Infrastructure Report (www.dhs.gov/dhs-daily-open-source-infrastructure-report)

INTRODUCTION TO PART III

PREVENTION: TECHNICAL DEFENSES

The threats and vulnerabilities described in Part II can be met in part by effective use of technical countermeasures.

The chapter titles and topics in this part include:

23. **Protecting the Physical Information Infrastructure.** Facilities security and emergency management
24. **Operating System Security.** Fundamentals of operating-systems security, including security kernels, privilege levels, access control lists, and memory partitions
25. **Local Area Networks.** Security for local area networks, including principles and platform-specific tools
26. **Gateway Security Devices.** Effective recommendations for implementing firewalls and proxy servers
27. **Intrusion Detection and Intrusion Prevention Devices.** Critical elements of security management for measuring attack frequencies outside and inside the perimeter and for reducing successful penetrations
28. **Identification and Authentication.** What one knows, what one has, what one is, and what one does
29. **Biometric Authentication.** Special focus on who one is and what one does as markers of identity
30. **E-Commerce and Web Server Safeguards.** Technological and legal measures underlying secure e-commerce and a systematic approach to developing and implementing security services
31. **Web Monitoring and Content Filtering.** Tools for security management within the perimeter
32. **Virtual Private Networks and Secure Remote Access.** Encrypted channels (virtual private networks) for secure communication, and approaches for safe remote access

33. **802.11 Wireless LAN Security.** Protecting increasingly pervasive wireless networks as a majority of the world's Internet users access information through their mobile phones and tablets
34. **Securing VoIP.** Security measures for Voice over IP telephony, a cost-effective method for increasing long-distance collaboration
35. **Securing P2P, IM, SMS, and Collaboration Tools.** Securing collaboration tools such as peer-to-peer networks, instant messaging, text messaging services, and other mechanisms to reduce physical travel, and to facilitate communications
36. **Securing Stored Data.** Managing encryption and efficient storage of stored data, including cloud-storage services
37. **PKI and Certificate Authorities.** Concepts, terminology, and applications of the Public Key Infrastructure for asymmetric encryption
38. **Writing Secure Code.** Guidelines for writing robust program code that includes few bugs and that can successfully resist deliberate attacks
39. **Software Development and Quality Assurance.** Using quality assurance and testing to underpin security in the development phase of programs
40. **Managing Software Patches and Vulnerabilities.** Rational deployment of software patches
41. **Antivirus Technology.** Methods for fighting malicious code
42. **Protecting Digital Rights: Technical Approaches.** Methods for safeguarding intellectual property, such as programs, music, and video, that must by its nature be shared to be useful

CHAPTER 23

PROTECTING THE PHYSICAL INFORMATION INFRASTRUCTURE

Franklin Platt

23.1 INTRODUCTION 23·2

23.2 SECURITY PLANNING AND MANAGEMENT 23·2

23.3 STRATEGIC PLANNING PROCESS 23·3

23.3.1 Attractive Targets 23·4
23.3.2 Defensive Strategies 23·4
23.3.3 Who Is Responsible? 23·5
23.3.4 One Process, One Language 23·5

23.4 ELEMENTS OF GOOD PROTECTION 23·5

23.4.1 Segmented Secrets 23·5
23.4.2 Confidential Design Details 23·6
23.4.3 Difficulties in Protecting the Infrastructure 23·7
23.4.4 Appearance of Good Security 23·7
23.4.5 Proper Labeling 23·8
23.4.6 Reliability and Redundancy 23·8
23.4.7 Proper Installation and Maintenance 23·9

23.5 OTHER CONSIDERATIONS 23·10

23.5.1 Equipment Cabinets 23·10
23.5.2 Good Housekeeping Practices 23·10
23.5.3 Overt, Covert, and Deceptive Protections 23·10

23.6 ACCESS CONTROL 23·12

23.6.1 Locks and Hardware 23·12
23.6.2 Card Entry Systems 23·13
23.6.3 Proximity and Touch Cards 23·13
23.6.4 Authentication 23·14
23.6.5 Integrated Access Systems 23·15
23.6.6 Portal Machines 23·16
23.6.7 Bypass Mechanism 23·16
23.6.8 Intrusion Alarms 23·17
23.6.9 Other Important Alarms 23·18

23.7 SURVEILLANCE SYSTEMS 23·18

23.7.1 Surveillance Cameras 23·19
23.7.2 Camera Locations and Mounts 23·20
23.7.3 Recording Systems 23·20
23.7.4 Camera Control Systems 23·20
23.7.5 Broadband Connections 23·21

23.8 FACILITIES DESIGN 23·21

23.8.1 Choosing Safe Sites 23·22
23.8.2 Physical Access 23·23

23.8.3 Protective Construction 23·24
23.8.4 Using Existing Premises Alarms 23·27
23.8.5 Clean Electrical Power 23·28
23.8.6 Emergency Power 23·30
23.8.7 Environmental Control 23·36
23.8.8 Smoke and Fire Protection 23·38

23.9 MITIGATING SPECIFIC THREATS 23·41

23.9.1 Preventing Wiretaps and Bugs 23·41
23.9.2 Remote Spying Devices 23·42
23.9.3 Bombs, Threats, Violence, and Attacks 23·42
23.9.4 Medical Emergencies 23·43

23.10 COMPLETING THE SECURITY PLANNING PROCESS 23·43

23.10.1 All-Hazard Mitigation Plan 23·43
23.10.2 Cost-Benefit Analysis 23·44
23.10.3 Security Response Plan 23·45
23.10.4 Implementation, Accountability, and Follow-Up 23·46

23.11 FURTHER READING 23·46

23.12 NOTES 23·47

23.1 INTRODUCTION. Chapter 22 in this *Handbook* describes a systematic approach to identifying threats to the physical information infrastructure. This chapter describes practical recommendations for implementing good physical-security protection for information systems (IS).

There was a time when insurance was enough to cover most threat situations. However, today disaster, workplace violence, and terrorism insurance coverage may not be available or affordable without an independent, outside security audit to ascertain a good security program is in place.

There is one nearly ironclad legal defense against negligence allegations, and this is to demonstrate that the organization is following the commonly accepted federal security planning and management procedures. Usually, such compliance alone is sufficient to prevent negligence suits from even being filed; any remaining allegations are usually minor and easily covered by insurance. Conformity to the commonly accepted federal procedures can quickly be ascertained by a recent security audit.

The security audit will examine compliance with all documented facets of the planning, preparation, policies, training and exercises, management and oversight, and response and recovery plans and procedures. These are all the necessary components for protecting the information infrastructure.

Chapter 54 in this *Handbook* discusses audits, and Chapter 60 discusses insurance for information systems.

23.2 SECURITY PLANNING AND MANAGEMENT. In the United States, there are many laws, regulations, and federal directives that should be incorporated into the security planning and management of any organization. Even when compliance is not yet mandated, it is highly recommended simply to avoid unnecessary and potentially huge costs of defending against allegations of negligence, and establishing in court that the security procedures in place are sufficient.

Some of the new requirements that may affect almost any organization in the United States are briefly summarized in the following list.

National Incident Management System (NIMS) Compliance: The National Incident Management System (NIMS) defined by the Federal Emergency

Management Agency (FEMA) requires all departments and agencies of the federal, state, and local governments to implement its standards. NIMS is an emergency-response template which supplements and unifies the incident command systems long in use by most response agencies.[1]

National Response Framework: The National Response Framework (NRF) provides an emergency operations plan which all state and local governments have long been required to have. The NRF now unifies the content and format for all such plans. Compliance also requires certification of senior officials.[2]

National Infrastructure Protection Plan: The National Infrastructure Protection Plan (NIPP) defines recommendations for the protection of critical infrastructure and key resources. It proposes a unifying structure that includes partnership initiatives, a risk-management framework, and information sharing.[3]

Robert T. Stafford Disaster Relief and Emergency Assistance Act of 1988 (Public Law 100-707): In combination with the Disaster Mitigation Act of 2000 (Public Law 106-390), every state and community must have a current and FEMA-approved *All-Hazard Mitigation Plan* and a current and approved *Emergency Operations Plan* in place.[4]

The Sarbanes-Oxley Act of 2002 (SOX): Public corporations must report all of their internal financial controls, which may include those relating to security management—including future events that can materially affect future earnings. Most of the threats identified in Chapter 22 can do this, so that the possible mitigated costs of each threat as determined at the end of this chapter may need to be included in the SOX reporting.[5]

The Health Insurance Portability and Accountability Act of 2002 (HIPAA): All organizations that receive and hold personally identifiable individual health information must keep those data.[6]

The Gramm-Leach-Bliley Act of 1999: Financial institutions, lenders, advisors, accountants, and businesses that process or receive financial information must protect those data against unauthorized disclosure and modification.[7]

Emergency Management Accreditation Program (EMAP): EMAP, "an independent nonprofit organization, is a standard-based voluntary assessment and peer review accreditation process for government programs responsible for coordinating prevention, mitigation, preparedness, response, and recovery activities for natural and human-caused disasters. Accreditation is based on compliance with collaboratively developed national standards, the *Emergency Management Standard* by EMAP. Accreditation is open to all U.S. states, territories, and local government emergency management programs. Anyone can subscribe to receive standards and guidance materials."[8]

23.3 STRATEGIC PLANNING PROCESS.

The planning process must think ahead strategically to determine the best possible security options to add maximum value and contribute to profits, enhance productivity and performance, and avoid intrusiveness. The process of identifying all possible threat situations and a statistical method to predict the likelihood of injury or damage and the current vulnerability of the enterprise to each individual threat is described in Chapter 22. This chapter builds on the planning process, suggests ways to better protect the infrastructure and thereby reduce its vulnerability, and describes how to implement the security planning and management process.

Some facets of good physical protection to consider in the planning process are outlined next. These suggest some critical elements and how to protect them, and some possible weak points in the information infrastructure. Section 23.11 completes the planning and mitigation process.

23.3.1 Attractive Targets. Much of this chapter deals with what amounts to physical hacking. As with its electronic counterpart, physical attacks cannot be predicted as to when, where, or how they will occur. Electronic hacking from within the organization, over its networks, or from anywhere in the world via the Internet is still the best way to break into most information systems. However, physical intrusion can be harder to detect and locate, and often is more damaging and costly. In addition, threats from employees or contractors who are bad actors with *authorization* to be unsupervised in their work with or near the physical infrastructure of information systems are a serious threat.

Chapters 12 and 13 in this *Handbook* discuss the psychology of computer criminals and dangerous insiders.

23.3.2 Defensive Strategies. There are many effective defensive strategies to consider. The planning process must evaluate each approach and how best to combine them strategically to maximize effectiveness and minimize costs. These defensive strategies are common:

- *Prevention* so that specific threats do not affect the enterprise.
- *Deterrence* so that specific threats are not likely to occur.
- *Mitigation* to reduce each threat to tolerable consequences.
- *Redundancies* so there are no critical links in the infrastructure that cannot be bypassed. There are many methods of redundancy, such as multiple data paths; bidirectional data loops; parallel or distributed processing; alternative support systems and utilities; and many more.
- *Early warning* to detect impending trouble and delay onset, so that fast response can prevent or minimize any disruption.
- *Layers of security,* which are like the concentric layers of an onion, so that several layers of security must be penetrated before a target can be reached. This adds reliability, because a failure or breach of one layer does not compromise the other concentric layers.
- *Insurance* that can reimburse some of the recovery costs but usually few of the response costs. Insurance coverage often excludes many threats and can be costly. Insurance may not cover (or even be available for) gross negligence, some acts of God, flooding, terrorism, or acts of war.
- *Capital markets,* which are less costly than insurance and better able to lay off larger and broader risks.
- *Self-insurance* to establish retentions (which are funds accumulated for the purpose) in the hope nothing serious ever happens.
- *Contract security services* that are performed in-house, which basically transfer risks, but do not necessarily mitigate threats.
- *Outsourcing,* which is another option that introduces still other and often unrealized threats and vulnerabilities.

23.3.3 Who Is Responsible? Effective security requires both governance from the boardroom and oversight by senior management. Top management must also actively sponsor it and insist that everyone involved understands, supports, and respects the security plans and procedures. Accountability and oversight are essential, as is insistence on periodic security exercises, review, and updating.

If possible, the infrastructure-security manager should *not* also be in charge of IS or corporate or premises security. Although managers of these areas face similar threats and their groups must work closely together, their priorities and levels of response to any incident are different. It is best therefore that two or more specialized security groups report to one senior executive officer with clear lines of authority. Even if the role cannot be full time, defining a formal position for an infrastructure-security manager is a wise investment.

23.3.4 One Process, One Language. As mentioned previously, the security planning and management process must be uniform and comprehensive, and applied effectively and efficiently. There should also be a current audit done by an outside, independent group to ascertain both security preparedness and compliance with the pertinent laws, codes, and regulations. The least costly and most effective approach is to follow federal guidelines. The Department of Homeland Security (DHS) and FEMA procedures provide a uniform and comprehensive methodology and generally accepted standard practices that are uniform, comprehensive, and most likely to deter liability as well. This process is recommended for any organization, public or private.

FEMA makes the distinction between an emergency, which is a situation that an organization can handle with its own resources, and a disaster, which is when internal resources are likely to be overwhelmed and outside help is needed. The distinction is not always clear-cut. "Crisis" is not a defined term. FEMA also recommends planning and preparation for worse-case situations, which often turn out to be what happens. Many threats can become disasters to a business, even when outside resources respond quickly. For example, the consequences of an equipment room fire can be disastrous when the local fire department responds with axes and water hoses, cuts off all electrical power to the building, and smashes out windows to vent smoke. Even a major incident that occurs away from the immediate premises can necessitate an evacuation, cutting off building power, and creating serious disruption.

23.4 ELEMENTS OF GOOD PROTECTION. Protecting the infrastructure generally requires different and stronger defenses than premises security or IS logical security can provide. The infrastructure protection must be effective, efficient, and affordable. Yet it must also be nonintrusive and user friendly. Too much protection is unnecessarily costly and often counterproductive. Too little can be even more costly and endanger productivity, morale, and goodwill.

Some of the requisite elements for just the right amount of protection are provided in the sections that follow.

23.4.1 Segmented Secrets. Although *security by* obscurity is rightly derided if it is the primary focus of protection, to maintain good security, no *one* person should know or control *all* the details or inner workings of the security systems and procedures. If total understanding of the security systems is segmented into several parts and the knowledge distributed among collaborators, there is much less likelihood of misuse, fraud, or error, and less dependency on a few key persons. The shared knowledge contributes to separation of duties and pervasive vigilance.

However, the more people with knowledge of the security systems and procedures, the more these systems become vulnerable. It is easier to compromise security through coercion or extortion, or by an unwise remark inadvertently disclosing information. The list of those with inside knowledge may include managers, administrators, maintenance personnel, users, partners, suppliers, customers, vendors, and consultants. Although many individuals must know at least some of the security protections, no one needs to know all of the details.

Beyond segmenting, there is another precaution needed for good protection. If members of the same group must share sensitive knowledge, a "two-person" rule should apply—an example of separation of duties. This principle is especially important when anyone is able to modify the security infrastructure, alarms, or event logs. The two-person rule says that any modification to the security system requires two authorized persons working together and that there will always be an audit trail showing who did what and when. This procedure also is used when two or more groups share responsibility for common elements of a security system.

23.4.2 Confidential Design Details. It is often quite easy for others to locate and identify critical infrastructure components. Many times the information is clearly shown on public documents. The signage on infrastructure locations often clearly describes what is inside, and critical components are often put in spaces easily accessible to outsiders. These are all unsafe design practices.

Many types of documents clearly indicate IS areas, support systems, and other infrastructure. The documents may show the locations of equipment, wiring, utilities, and cable runs. Documents that are likely to reveal sensitive information include building plans that show floor and office layouts, furniture, wiring, and equipment locations. Architectural and engineering documents, drawings, and specifications often show sensitive information. These documents often list the function of each area, and even an occupant's name or title. Other types of documents breach security as well. Examples may be as-built plans, alteration and construction plans, electrical and communications wiring diagrams, patch-panel and cross-connect setups, as well as contract drawings or proposals, shop drawings, installation plans, maintenance diagrams, and, especially, documents filed with code-compliance and regulatory agencies. Good protection requires that all such documents be controlled and kept securely stored. Better yet, sensitive information should be removed and alphanumeric designations that are cross-referenced to sheets than can easily be kept classified should be shown.

All of the listed documents are routinely distributed to a wide range of sources, such as interested contractors and bidders, vendors, suppliers, and maintenance providers. Building managers, landlords, and often real estate offices are likely to keep copies on file, and many other persons can readily access the documents and copy them. Most such documents are publicly available or obtainable via the sunshine laws of each state, by court order, or simply by deception. Moreover, many legitimate persons obtaining or receiving the documents have no internal document control or security provisions.

All room and area designations should always be alphanumerical. Descriptive or functional names should not be used for any area. This caveat pertains to the entire premises, all public areas of the facility, and all building mechanical and core areas. Functional designations or terms such as "treasurer," "marketing director," "security desk," "computer room," "network closet," or "telephone room" should never appear. Nor should the names of any occupants, departments or functional groups, or individual tenants' spaces be given. Use only alphanumeric designations.

Lists that correlate area and room numbers, functional areas, or any descriptive names should be kept secure in a locked file, as should equipment room drawings, patch panel connections, and wiring plans. All these must be readily available to system administrators in the event of a system failure or when doing maintenance or upgrades, but only on a need-to-know basis. Ideally, even managers do not have access to this information, except when accompanied by security personnel. As well as controlling access to such documents, it is equally vital that every document be kept current.

Finally, security personnel should review all plans, drawings, and documents for construction, alterations, equipment moves, or any other physical changes before the information is issued. Security personnel should also review the invitations to bid and all the drawings and specifications, review again the quotations received, and continue to review every as-built document produced throughout a project. Internal security should review all these as a matter of policy. If the security personnel do not have sufficient time or expertise to review everything quickly and thoroughly, outside independent experts can be valuable. When properly chosen, these experts can provide broad experience, perspective, and evidence of due diligence, so that management cannot later be accused of negligence in protecting information, people, or property. Finally, the security leader should sign-off that no project documents violate any internal security policies.

23.4.3 Difficulties in Protecting the Infrastructure. Decades ago, most IS equipment was housed inside a single computer room, and most of the external terminals and peripheral equipment were located nearby. Cable runs were short and generally within secure areas, and access controls, alarms, and surveillance could easily cover the critical areas. There were often security guards as well. But in those days, there were fewer threat situations and less cost if trouble did come. Many organizations then were apathetic and security was lax. They were lucky, but a few incidents did occur and often at great cost. There were also likely many undiscovered security breaches and still more that were never reported.

Today, the IS infrastructure is much larger and more complex, and the potential costs of trouble are far greater. Today's infrastructure is much more interdependent, and it now includes many more equipment and network rooms. It extends to many more telephone and utility closets, and interconnects widespread and diverse desktop, peripheral, and remote nodes. Today's infrastructure is increasingly harder to protect, and the future outlook suggests many more threat situations with the potential to cause major business disruption.

To further complicate security, there are now many more and diverse IS interfaces and a complex infrastructure to protect. Interfaces now include direct and switched wiring, wireless topologies, and infrared coupling. Access to the Internet, local area networks, and wide area networks may now utilize combinations of metal wiring and fiber optics, wireless, satellite, TV cable, and microwave links. Some of these interfaces will be dedicated, others switched, and still others temporarily patched. It will be difficult to even locate all the interfaces, yet each must be protected, as must be their cable runs and the utilities that support them. Today, the early warnings from the security alarms and defenses must be so effective that most trouble can be prevented before it happens.

23.4.4 Appearance of Good Security. The appearance of an armed fortress is usually counterproductive. This usually intimidates and obstructs both visitors and staff more than it protects them. The same is often true for too many guards and

receptionists stationed to block entry points to internal areas (unless this is deemed absolutely necessary). Such barriers tend to enrage anyone who may be already anxious, and can provoke attacks. Barriers tend to be ineffective; for example, even if the glass is bulletproof and spray-proof, the pass-through holes likely are not. Most barriers can be breached as well, and the person(s) inside injured or at least traumatized. For example, some years ago most banks rushed to put their tellers behind thick glass panels. However, they soon discovered that the glass provoked trouble rather than deterring it. The banks quickly removed both the glass and the pass-throughs, which many businesses then bought and installed, with similar results. Anger and potential violence are best avoided by good facilities design, security systems and access controls, and training in security awareness and violence prevention.

There are also considerations whether security devices should be covert or appear in plain sight for all to see. Defenses that are in plain sight can serve as deterrents and promote a feeling of safety, but they can also be vulnerable themselves. Cameras can be spray-painted, shot out, or knocked aside with a club. Any exposed wiring (which is a security no-no in itself) can be cut or shot through. The alternative is to use concealed devices and also dummy devices or "honeypots" that are obviously positioned. (Both are discussed later.)

23.4.5 Proper Labeling. Good security requires quickly locating trouble spots, with certainty that the nomenclature of every component, cable, and connector is clear and current. There must be proper and consistent labeling of all data and power cables, both inside and beyond the secure areas. No label or tagging should reveal confidential information. Many vendors, installers, contract personnel, and in-house staff tend to use their own labeling and tagging systems, some of which are clear and understandable, while others are not. The labeling and tagging must match the documentation, plans, and drawings as well. Sloppy wiring management is often the norm as changes are made but not marked or documented, or wiring abandoned and not removed. The status of in-house personnel, installers, and maintainers is likely to change frequently, creating all the more potential for misinformation and confusion. The inevitable result is poor protection and slow response to a security incident. There are generally accepted wire-management procedures, and a single method should be utilized throughout the information infrastructure. It is especially important that all authorized personnel understand and accept the labeling system.

The Telecommunications Industry Association/Electronic Industry Alliance (TIA/EIA) publishes the generally accepted authority, TIA/EIA Standard 606-B, *Administration Standard for the Telecommunications Infrastructure*,[9] which describes labeling of cables, connectors, patch panels, cross-connections, and equipment. This standard also requires labeling firestops, grounds, special-ground circuits, and neutral wires. The *National Electric Code* (NFPA 70),[10] published by the National Fire Protection Association, also includes standards for labeling cables and conduits. Local codes may augment the national codes. Generally, any substantial alterations, moves, or changes, and usually all new construction will require full compliance with the current standards throughout the premises—which is a wise security investment as well.

23.4.6 Reliability and Redundancy. The first requisite of reliable system performance is reliable equipment, systems, and infrastructure, properly installed and maintained. Unfortunately, some components may be poorly designed, subject to erratic quality control, damaged in transit, or applied wrong. Much equipment now includes a fail-over mode to maintain full performance when a failure occurs.

Another component of high service levels is *redundancy*: concurrently accessible equipment and parallel paths to take the load should one component falter. Effective redundancy for business continuity also requires alternate sites that are off premises to process and store information. It is best that there be multiple alternate sites, each within a safe environment and well distant so that problems at the primary site will not affect any alternates.

Inside each computer room, storage systems should be redundant arrays of independent disks (RAID) allowing important data to be fully replicated. For critical data, the RAID storage systems should themselves be redundant. Servers should also be redundant. Multiple parallel servers with load balancing are a wise investment, so that one server will automatically take over another's load if it falters or must go off-line for any reason. The load-balancing features also provide potentially higher throughput for the overall system by allowing parallel computations and input/output. All critical equipment should have redundant power supplies, fans, and hard drives that can be diagnosed quickly and hot-swapped easily.

Good system manageability is another vital requirement. This includes hardware that can detect trouble before it happens and, if possible, pinpoint the exact location. It also includes good management software with warning and alert capabilities and good logging systems. Remote management capabilities must also be private and secure to preclude penetration or denial-of-service attacks. Good security management also requires that any changes or disabling of alarm parameters should require two authorized persons to be physically present inside the equipment room and simultaneously logged on. Good system management adds some cost but is a wise investment in effective security and oversight.

23.4.7 Proper Installation and Maintenance. Good protection requires that all information systems, equipment, and wiring be installed properly, according to the manufacturer's instructions and the intended usage. All the wiring must conform to, or exceed, local code requirements. Data wiring should be independently tested and certified that it meets specifications, current standards, and, if possible, anticipated future needs.

Out-of-the-box equipment hookups and installations are common and the cause of many system failures. Most security features are disabled when components and software are shipped. Proper installation requires careful setup, customization, and performance testing, for which adequate time and resources must be allocated. Promptly installing the latest modifications, service packs, updates, and security patches is also vital to maintaining performance. Delays of days, weeks, and even months often intervene, while new threats emerge or threat levels escalate until the new defenses are in place. Once installed, the information systems and infrastructure must be periodically reviewed, tested, and kept up to date. Frequencies of such reviews and testing must reflect the known or estimated rates of failure and the consequences of such failures.

Administrators, installation, and maintenance personnel must be properly trained, experienced, and, in many cases, certified or licensed as well. Each person's credentials should be checked before being permitted on site. Given the limited staff, time, and budgets available, there is often more lip service than actual certainty in the process of reliability assurance. Management must understand that proper installation, upkeep, and maintenance together constitute cheap and effective assurance that IS performance is never compromised.

For more details of data backup methodologies, see Chapter 57 in this *Handbook*; for business-continuity planning and disaster recovery, see Chapters 58 and 59.

23.5 OTHER CONSIDERATIONS. There are many other factors that can significantly establish and improve physical security. A few of these are discussed in the next sections.

23.5.1 Equipment Cabinets. Most IS equipment is now open-rack mounted to save floor space, for more reliable performance, and for easier access. Although enclosed equipment cabinets cost more than racks, they offer much better protection. Equipment mounted within a closed cabinet can be better ventilated and cooled, kept freer of contaminates, and may also escape damage from external particulates, water, smoke, or liquids. Should there ever be overheating or smoke generated within an enclosed cabinet, the condition can be detected quickly and usually resolved before equipment or wiring are damaged. Locks keep out those unauthorized, or at least delay access and leave evidence of trouble. Open-door alarms provide another strong layer of protection, and wiring conduits to cabinets can also be protected. Cabinets with redundant fans can better monitor and maintain ventilation and cooling, which, in turn, facilitates more equipment mounted in less floor space. In all, closed cabinets can provide an additional layer of protection against accidental or deliberate damage to the infrastructure.

23.5.2 Good Housekeeping Practices. All food and drink must be kept out of equipment rooms, since they can cause considerable damage if spilled on equipment, a monitor, keyboard, connectors, and wiring and cable harnesses. Food also attracts insects and rodents (which are found in many buildings), many of which also like to eat wiring. In addition, grease and other contaminants from food can easily be spread by operators onto critical surfaces of equipment such as input output devices and cause read/write errors or even complete shutdown of the devices. For example, a single greasy DVD can ruin a DVD unit or high-capacity backup-tape unit. Worse, a contaminated read/write head or entry slot can in turn spread contamination to other media and then on to other devices. Space for food and drink should be provided outside the equipment room, where routine maintenance personnel can keep the food area clean. A convenient dispenser for alcohol-based hand wipes can encourage cleanliness for everyone entering the equipment room.

Loose papers, books, supplies, newspapers, and trash are fire hazards and also must be banned from every equipment room.

23.5.3 Overt, Covert, and Deceptive Protections. Effective protection of the IS infrastructure requires many hidden elements—such as concealed surveillance cameras, sensors, and detectors—and all of the wiring that supports them. But good protection also needs some clearly visible elements. It is important to consider which devices are best hidden to protect them and which should be visible as deterrents.

Overt devices are ones that are evident to workers and visitors, or whose presence is implied by other visible objects, such as warning notices. These visible devices, which suggest that some sort of security exists, are intended to deter troublemakers, so that all but the most determined attackers will go elsewhere. Examples are surveillance cameras, access controls, visible alarm boxes, and visible sensors. Although most overt devices are active and recording data, some may be inexpensive dummy devices that only look real, perhaps with slowly blinking indicator lamps to heighten the effect. Covert protection, however, must not be noticeable to either visitors or insiders. There

must be no indication that these protections exist, what they are, how they might function, or where they are located. Most effective security systems operate covertly; examples include stealth and silent alarms, concealed early-warning systems, perimeter and proximity sensors, access monitors, and many other surveillance devices that are not readily seen.

It is important also to conceal the wiring that interconnects all protective systems and the utilities that support them. Whether any part of a system is visible or not, the wiring that connects it should not be. Although overt devices may themselves be vulnerable, they will generally advertise that there is good security here and everyone within the premises can feel safe. However, visible devices can sometimes be covered or spray-painted, knocked aside with a club, or shot to disable them. An expert thief can also defeat many hidden systems, if he or she knows what they are, where they are located, or how they are connected.

Another approach to protection involves deception. Dummy devices that look like surveillance cameras, access control devices, and alarm sensors can be placed to attract troublemakers, who may think they can physically damage, disable, or circumvent the system. These visible devices are intended to distract potential troublemakers and divert them away from vulnerable areas. Some devices are deceptive in that they are not what they appear to be but are actually alarm sensors to measure motion, proximity, sound, or anything that disturbs the device. Deceptive devices often are used to divert troublemakers away from vulnerable people and infrastructure, by offering them a "honeypot": an attractive target to distract them, but often a target equipped with an alarm device, surveillance cameras, or other means of identifying a perpetrator and gathering evidence.

There is a gray area between what management can legitimately do to protect its information systems and what may be unethical or illegal actions. Management has a legal and fiduciary responsibility to protect people and property, and those who support deception say that these techniques are increasingly necessary to protect an organization. Others insist that this amounts to entrapment or violates privacy rights. State and local regulations and interpretations vary widely and are continually changing. It is necessary to check carefully with local officials, legal advisors, and insurers to determine what is acceptable and how to manage such risks. Management must then decide to what extent these techniques may be effective and whether less contentious approaches will suffice.

Whether the protection devices themselves are overt, covert, or deceptive, the security systems behind them must not be obvious. No one seeing or knowing about the elements of a security system should be able to deduce the details of the system, the functionality, or where and how it is monitored. An observer may notice a particular device or product or the suggestion of a vendor's standard security solution, but the particulars of the protection systems must remain obscure, and all the wiring that supports them must be hidden or disguised as well.

Everyone involved must be aware of the security policies and procedures. Conspicuous signs should advise that anyone entering the premises may be monitored, as may all communications. All of the security policies and procedures should be understood and accepted by everyone involved. Employees and other on-site personnel should receive periodic security-awareness training and briefings. And there should be periodic security exercises and drills to test the procedures and reinforce the training.

Finally, protection must not be intrusive. Security cannot limit productivity or IS performance in any way. Instead, the protection must contribute to a feeling of safety and security within the workplace and thereby enhance productivity.

23.6 ACCESS CONTROL. Access control systems are but one layer of good infrastructure protection. They are usually used in conjunction with surveillance and perimeter control systems in order to provide sufficient early warning to head off trouble before it happens. Effective access control requires three tiers of support that are described next. The strength of each tier and its integration with other security layers determines the security effectiveness.

1. **Privileges.** This tier determines whether a person seeking entry is authorized. It is the initial entry-request process that may use an access or proximity card, radio-frequency identification (RFID) key, or a combination lock. Since many of these devices can be lost, borrowed, stolen, or copied, and many can be quickly defeated, there is usually not much effort to ascertain just who is seeking entry. Therefore, privileges alone are not strong security.
2. **Authentication.** It is usually necessary to identify a person seeking entry with some degree of certainty. To do this, the person must possess or know something unique that others cannot readily duplicate. Examples include personal identification numbers (PINs), electronic keys, entry cards, and biometric devices. PINs and passwords may be used, provided they are strong and well implemented. Some of these approaches merely strengthen the privileges process but can still be copied or defeated.
3. **Audit trail.** A log is required for each entry attempt to show the date and time, the identification of the person, and the action taken by the access control system. Access-denied and unable-to-identify events should trigger immediate alarms. Logs must be analyzed in a timely manner for anomalies or unusual patterns. Where better access control is needed, each person's exit also must be authenticated and logged.

See Chapter 28 for more details about identification and authentication in access controls.

23.6.1 Locks and Hardware. Strong protection begins with high-quality locks, door hardware, and access control systems that are nonintrusive yet strong enough to deter most unauthorized entry. Lock types should be hard to pick and should use keys that are hard to duplicate. Examples are Medeco® locks and keys with dimpled sides. Ace® locks with circular keyways require special tools to pick. Many types of keys can be created from lock numbers, so keep these numbers stored securely.

No lock is completely safe. Someone with equipment, experience, and time can open any lock, often very quickly and without causing attention. Where key mastering is used, an experienced person with a key to any single door can open the lock cylinder and copy the mastering system. Therefore, additional layers of protection are needed.

Interior areas accessed only by a few people usually can be secured with a strong push-button combination lock. Key locks are not appropriate, because the keying cannot be changed periodically or quickly when a key is lost or someone leaves. And misplaced or lost keys may not be reported for days or weeks. However, key locks tend to be stronger and less vulnerable to vandalism, so keys may be the best alternative for outside areas or for doors that remain open during business hours. Wherever keys access critical areas, there should be spare lock cylinders and a new set of keys stored on site that can be utilized quickly when a change is needed. Once a lock cylinder is changed, the old cylinder should be rekeyed and a new set of keys produced.

Locks that use an electronic key are particularly effective. Electronically activated cylinders can replace existing mechanical cylinders, and many do not require any wiring, so the hardware and operating costs are minimal. Most electronic keys have a small cylindrical tip that is touched to the lock for access. Both the lock and the key can log each event, identify the specific lock or key used, the date and time, and whether access was granted. And conveniently, electronic keys are not much larger than mechanical keys. However, both can be defeated with the proper equipment and skill. Often, two independent entry locks are utilized to provide stronger, relatively inexpensive protection.

RFID devices can provide good access control, provided the activating device is not lost or stolen. RFID can be an improvement over card entry systems in that it can activate a lock from several feet away. RFID access cards can include photo ID, name, and perhaps an optical stripe for encrypted identity information. As with other systems, RFID can be breached with the right equipment and skills.

Another inexpensive upgrade of key locks is the card-access lock similar to the ones used by hotels. Many do not require wiring and are battery operated, so the keying of each door remains unchanged and an old key remains valid. Therefore, without central wiring, most card-access locks offer limited security and cannot trigger an alarm although some do at least log events.

To prevent intruders from bypassing the lock mechanisms by using credit cards or thin metal sheets to push the bolts inward, the area over the actual latch should be protected with an *astragal*—a flap of metal bolted to the door. In addition, latches should include anti-tampering features that block external movement of the moveable elements into the door when the locks are active. The hinges for doors should not be accessible from outside the secured area when the doors are closed: It is a simple matter to break and remove the hinge bolts and remove a locked door if the hinges are on the outside.

No matter how strong the access control systems, doors have their limitations. Absent a vault-type door and hardware, a determined attacker with pistol and silencer can gain access readily. A small water cannon that is transported in what looks like a toolbox can breach a standard metal door, or the partition surrounding it, with one shot and very little noise.

23.6.2 Card Entry Systems. One of the best means of physical access control is a card entry system, especially the newer systems that are increasingly capable and less costly as well. A central card entry system often controls the entire premises: all entrances, exits, elevators, rest rooms, and many other interior doors. Access cards are usually similar to a credit card and can be carried concealed or worn as identification badges. Access cards are usually imprinted with a full-face photo, the individual's name, the organization's logo, and often a printed stripe with biometric information. Access cards often show number codes to indicate authorized areas of access, and are usually color-coded to indicate status and whether the person must be escorted.

For more details of identification and authentication systems and card entry systems, see Chapter 28 in this *Handbook*.

23.6.3 Proximity and Touch Cards. Some cards can be held near a sensor (*proximity cards*) and some can be held in contact with a sensor (*touch cards*). Both types of tokens can reduce wear on physical tokens that have to be put directly in contact with the sensors.

For more details of identification and authentication systems and proximity and touch cards, see Chapter 28 in this *Handbook*.

23.6.4 Authentication. Anyone can use an access control token that may be borrowed, lost, copied, altered, stolen, or taken by force. Therefore, authentication is another layer of security that is needed to establish the identity of the person seeking entry with some degree of certainty. Authentication devices commonly include a biometric scanner, a numeric keypad, and visual or voice identification by a computer or by another person. All such devices come in varying security strengths and each can be used in combination.

When an access token is read, the system must verify that the *token* has the requisite privileges for that place, date, and time. If it does, it then becomes necessary to authenticate the identity of the person using the token. For this purpose, a numeric keypad was once the most common device whereby the token user entered a personal identification on a number keypad. Now most use a touch screen where the numbers will appear in random order. If the system validated the PIN, it activated an electric door strike to momentarily unlock the door. This system can be slow, cumbersome, and easily compromised. In some systems, everyone uses the same PIN and is supposed to keep it private. However, PINs can be forgotten, lost, or discovered by others.

Visual authentication is a better approach. This can be done by a computer or by a security guard or receptionist who can see the entrance or monitor it via a surveillance camera or a video intercom. Emergency-type video intercoms work well because they provide the remote authentication with a visual image and can monitor sound continuously, so that the authenticator can speak with and challenge whoever approaches the door. When identify is verified, the system or other person activates the electric door strike to unlock the door. This system offers stronger security and is faster than a keypad. It also facilitates recording and logging all entrances and exits, especially during off-hours. Breaches can happen, if the person seeking entry can deceive the guard or receptionist.

Biometric scanners offer the strongest security. They offer faster, more positive identification of every individual and do away with the need to remember a PIN. Biometric scanners can read any of these personal attributes, which are listed somewhat in order of their current popularity: fingerprint patterns, facial characteristics, voiceprint, handprint, signature, retinal details, or height and weight.

Most biometric systems can be adjusted to be highly sensitive (which is slower and may require repeated entry attempts, but is very hard to breach) or less sensitive (which is still fast but may result in some false authentications). Before choosing a biometric system, it is necessary to determine that the users understand and will accept it. For example, some people balk at having to use a retinal scanner, while others may feel that any biometric device invades their privacy. The latter reason is invalid because most biometric systems cannot be used to identify any person not already known to the system. Increasingly, especially in critical public places, these systems can check every individual against criminal and terrorist databases and quickly alert law enforcement. It is difficult to steal an identity using biometrics because one must also know the encoding algorithms used.

Fingerprint scanners are the most common and are becoming increasingly powerful. Initially, these utilized optical scanning, which could be fooled by photographs, wax impressions, or by a severed digit. The newer capacitive scanners use electronics rather than optics and can provide nearly certain identification. Usually, three of the user's fingers are "enrolled" in the system in case some fingers are later bandaged or dirty. If

all the enrolled fingers are incapacitated, single-use passwords can be used to bypass the system.

Most fingerprint scanners cannot identify an individual by name, but only that a person seeking entry matches the person whose biometric identity has been enrolled. Most scanners do not conform to the uniform, automatic fingerprint identification standards used by law enforcement. Instead, they scan a small area of the finger and apply proprietary algorithms and encryption. A template from one system is usually meaningless to another.

Accuracy of fingerprint scanning is affected by the angle and position of the finger and by the pressure applied. Most systems allow sensitivity adjustment to optimize enrollment and verification times and success rates as well as to minimize delays and false negatives that require repeated access attempts. While well suited to most applications, fingerprint scanners may not be appropriate where the user could have dirty or thickly callused hands or must wear gloves (such as healthcare workers).

Facial recognition is the basis of another popular scanning system. It uses graphics technologies and any surveillance camera to measure the size and relationship of prominent facial elements. Most systems are proprietary and cannot be used to identify an individual by name, but some others are compatible with law enforcement standards. Sensitivity and accuracy are dependent on the distance, position, and angle of the head as well as on the background lighting. Cameras at all entry points must be positioned to photograph the subject at the same angle that their faces were enrolled. Not all facial recognition systems offer strong security; some can be fooled by a face mask or a photograph. Others are best left at their highest sensitivity settings to avoid spoofing, which may require increased false positives and multiple attempts at entry.

Voice-print scanning can be used, but mostly for access to a terminal or workstation. The better systems display random words on the monitor so a prerecorded response cannot be used. Most use the workstation's microphone or, increasingly, a cell phone. Voice scanners can be affected by hoarseness, so there should be a one-time password access provision. These systems can be useful for remote login, especially while traveling, although the low bandwidth of dial-up circuits may lessen the system's usefulness.

Retinal or iris scanners are considered the best security, but authentication can take a few seconds. For access control, the user must generally look closely into an eyepiece, which is traumatic to some people. But for access to a terminal or workstation, a Web camera is generally placed on top of the monitor, 17 to 19 inches from the eye. The user's head must be positioned to align a white dot within a colored circle. The user must hold still, without blinking, while the scan proceeds automatically. The camera can be used for surveillance to see who is near the workstation.

For details of biometric authentication, see Chapter 29 in this *Handbook*.

23.6.5 Integrated Access Systems. Biometric scanners are used increasingly by other applications, and most can readily coordinate with an access control system. Applications include network user authorization, and access to terminals and workstations and to software applications and data. Biometric readers can be built into laptops, keyboards, mice, or peripheral readers. They can be used with access cards, badges, and proximity or RFID devices for authentication with some degree of certainty. Scanners that are mostly proprietary are currently integrated with encryption systems (e.g., virtual public networks, public key infrastructure, and smart cards) to authenticate transactions including credit cards, financial, banking, and automatic teller machines. Biometrics are increasingly used to identify hospital patients, welfare

recipients, people who frequently enter the United States, and similar applications that involve identifying a diverse or widely dispersed group.

Infrastructure security protection can be independent of, or integrated into, a comprehensive premises security system. Either cards or badges can provide many other functions beyond basic access control. Applications include off-hour access to the building, elevators, and premises; control of building entrances, restrooms, and parking areas (especially at night); purchases at a cafeteria or company store, or for charging stationery or materials picked up from a supply room. They can also be used to receive classified documents.

For greater protection and efficiency, an integrated enterprise-wide, Web-based system can control access to various premises, locked doors, infrastructure components, networks, workstations, applications, secure data, and Web connections. This arrangement offers comprehensive security by logging every event. Centralized logs can yield much more meaningful security information, because an integrated system provides better early warnings to head off trouble before it happens.

23.6.6 Portal Machines. Airportlike security checkpoints are increasingly being installed outside airports. A portal machine is the archway one must walk through during the airport security process. It detects concealed metal, such as guns and knives, and tools that might be used to cause trouble. A portal and can be used to detect IS components (such as storage media) being smuggled into or out of the premises. For example, the Texas Capitol building had such a system installed in May 2010; however, Texans carrying concealed handguns need merely show their license to be able to circumvent the scans entirely.[11]

Newer trace-portal machines can detect many explosives as well. The person entering is asked to stand still for a few seconds while the machine releases several puffs of air and captures samples that are then analyzed for a number of hazardous or explosive substances. When a guard is not present, a computerized voice instructs the person to proceed, or a turnstile can hold the person in place.

Such devices might be appropriate for all who enter premises requiring ultra-high security and at entrances to critical areas or equipment rooms as well.

In January 2013, the Transportation Security Administration (TSA) announced that 200 specific full-body scanners would be removed from U.S. airports because of concerns over the resolution of the nude images produced by the Rapiscan™ machines.

In December 2012, the DHS announced that it would contract with National Academy of Sciences for a thorough study of the intensity of electromagnetic radiation in security scanners.[12]

23.6.7 Bypass Mechanism. Whatever the systems used, there should be one bypass key that can open every door that uses electronic or electrical controls. The bypass key is for emergency use only. It is a unique key that is not on any mastering system and is available only to a few highly trusted people. The cylinders should be heavy duty and very hard to defeat, with the keys nearly impossible to copy. Careful control and protection of each key is essential. The loss of a key may not be discovered quickly, and the time and costs of rekeying every lock will be substantial. Meanwhile, every door on the access control system is vulnerable.

Bypass passwords for individual users also may be needed. These passwords should trigger a silent alarm whenever used, so that security can respond to the site or verify identity by telephone, surveillance, or intercom. One-time passwords provide the best security.

23.6.8 Intrusion Alarms. Intrusion alarms are necessary to provide perimeter and early-warning alarms and are usually needed as extra layers of security. There are several methods of intrusion detection. Digital surveillance cameras with motion detection are best because they can monitor visually what is happening and record what they see. Other methods include proximity and pressure sensors mounted within the perimeter walls or floors or inconspicuously within the room. Most of these sensors can detect intrusion and forced entry and can pinpoint the location of trouble, but they provide no details, monitoring, or evidence-recording capabilities. Proximity and pressure sensors can protect long-distance perimeters, cable runs, and utilities inexpensively. Concentric layers of such devices are necessary for sufficient early warning to prevent trouble from happening.

The best motion detectors use digital closed circuit television (CCTV) surveillance cameras that can sense movement while observing and recording the event. Miniature cameras that are inconspicuous or concealed are particularly effective and increasingly inexpensive. Several cameras can record pictures from many angles and often can identify an intruder positively. In a larger area, cameras often use swivel-tilt mounts and zoom lenses, which can be controlled automatically and remotely as well. Color cameras are preferable, as are cameras that automatically adjust to light conditions, including near darkness.

Some CCTV cameras include a military-type night-scope mode, which is relatively inexpensive and functions well in near-total darkness. These cameras also work well in normal light and can switch automatically to night-scope mode when necessary to see clearly and to record evidence. Other types of cameras, such as infrared, work well where there is no (humanly) visible light.

Other intrusion detectors use radar technology or measure changes in capacitance or inductance to sense intrusion. Most cameras and detectors can trigger an alarm as soon as they are disabled or lose power. Detectors may be wall- or ceiling-mounted devices as an overt means of deterrence, but these may then be vulnerable to spray paint, wire cutters, a gun, or a club. Therefore, intrusion detection sensors are usually concealed, or at least inconspicuous.

Perimeter alarms are especially important in building core areas, or public areas within or outside a building, to provide ample early warning of an intrusion that might soon affect a secure area. Digital video cameras are best for this but may be ineffective over large areas or long distances. Therefore, many proximity devices are used to monitor intrusion. Most utilize long sensor wires that are surface-mounted inconspicuously or hidden within partitions, ceilings, and sometimes inside of conduit. These systems detect the presence of a standing adult or a large animal that may come within a few feet of the sensor wire. The sensor wires can be very long, so zoning is often necessary to pinpoint an incident at least within a general location.

A better and cheaper alternative can be fiber optic perimeter alarms. Developed for the military and national security installations, the fiber optic systems are very sensitive and can monitor, evaluate, and record events. The sensor wires can also be embedded inside drywall or masonry partitions, ceilings, floors, or conduit and will detect both pressure changes and ambient sound. Because they do not measure proximity and can monitor and evaluate events, false alarms are less likely. They can warn of an impending accident or efforts at forced entry, and may soon be able to locate the event as well. These systems use software that can discriminate between recognized events and situations that are unknown or potentially dangerous. All the events are recorded and can be replayed and reviewed by a remote operator at any time.

Increasingly, inexpensive monitors are available that connect wirelessly via Wi-Fi, mesh networks, cell phone, satellite, a radio channel, or sometimes via microwave. Such devices save the high costs of direct wiring to remote areas and the need to protect and conceal the wiring. But there is also a trade-off with reliability and security, as any radio-frequency device can be monitored or jammed and may also be spoofed.

In addition to intrusion alarms, environmental alarms should be provided to measure temperature, humidity, and smoke, fire, and flood situations within all critical areas and equipment cabinets as well, as described in Section 23.8.1.

23.6.9 Other Important Alarms. A duress alarm system is recommended within most critical areas. This is usually a large red button or an emergency intercom conspicuously mounted near the exit. It is used if someone is injured or ill, when trouble occurs that other alarms will not immediately sense, or if violence is threatened. The emergency intercom types are best because each party can talk and listen. Those with a CCTV camera are described here, but many inexpensive emergency intercoms that provide audio only can be useful. Security personnel can constantly monitor all sounds within a secure area, and anyone inside can readily talk with security personnel. Duress alarms are usually silent. And activation devices can be concealed or located inconspicuously in case a potentially violent situation erupts. Duress alarm activators inside of cash drawers, in the knee wells of desks, or under counters have been common for years.

Another form of duress alarm can be integrated into card-access systems or password systems, for example, swiping the card backward may allow access but also silently trigger alarms at security stations. Proximity cards could be programmed to recognize peculiar methods of waving the cards which might pass inspection of those applying duress yet initiate appropriate defenses and alarms. Entering a keycode with an additional sequence could do the same. All such duress-alarm signals should be part of every employee's training for use of the security systems.

Beyond access control and authentication, it is also important to know whenever a locked door is not fully closed. There should be a sensor that warns whenever an access-controlled door is open or ajar. The sensor is normally built into the door buck (frame) to provide a silent alarm. An open-door alarm system delays for a few seconds in order for one authorized person to enter or exit normally. A short delay prevents leaving a door ajar for someone else to push open, piggybacking when another person enters or leaves, and tends to prevent anything large being taken in or out. The open-door alarm also prevents a door from being propped open at any time. The door-ajar sensors should be concealed at all times so users are unaware of their existence. Otherwise, they may be taped, jammed, or otherwise defeated.

A final security protection is to prevent *double-dipping* whereby an authorized person requests multiple entries in order to admit unauthorized persons. The access control system can prevent this; once persons enter a space, they must be logged out before they can try again to enter.

23.7 SURVEILLANCE SYSTEMS. Today, surveillance systems are designed and laid out to document fully every event, to facilitate positive personal identification, and to provide legal evidence when needed. Today's digital cameras, controllers, and recorders can do all this and more. The old analog or film cameras and recorders are inadequate and should be replaced, especially since the new surveillance systems are much less expensive.

If the protection is designed well, cameras can provide an undisputable, accurate, historical record that is available instantly. Cameras never sleep and are not distracted by repetitive tasks or boredom. More important, cameras can provide early warning, and can document events from many perspectives, concurrently and synchronously. Cameras increasingly incorporate microphones and interface with emergency audio intercoms, the better to assess a situation, assist people on the scene, or challenge suspicious persons. The very presence of visible but inaccessible cameras usually deters most troublemakers. Protecting the infrastructure requires both early warning and identifying the nature of trouble before it can happen. The new surveillance systems can do this effectively. And they can be integrated with other alarms and with premises security systems for strong, seamless protection.

23.7.1 Surveillance Cameras. Surveillance cameras are far more effective and much less expensive than guards, watchmen, or extra receptionists. However, for strong, redundant, and flexible security protection at important locations, both cameras and people will be needed.

Surveillance capabilities are becoming increasingly effective protective devices. They can now see more, detect some dangerous hidden materials, and, increasingly, decide what is important to monitor.

Equipment rooms and other sensitive areas once relied on motion detectors for security, but now digital cameras with motion detection work better. Digital cameras also function well indoors and outside, and under most light conditions, from sunlight shining into the lens to near, if not total, darkness. When used outdoors, digital cameras are immune to all but very severe weather conditions. They are, however, affected by heavy smoke, snow, ice, and rain. Wind-blown objects and sometimes birds or small animals can distract their motion-detection systems.

Today's digital systems are far more capable and reliable than the earlier analog or film cameras. Cameras were once the weak link because the details were not clear or too small. Many of the old cameras were needed to cover an area, and even then, necessary details were often out of focus. Now cameras are much smaller, less expensive, and take advantage of auto-focus, tilt, and zoom capabilities to yield much better pictures. Almost all cameras now use color as well for better clarity. Some also incorporate an inexpensive, monochrome, night-scope mode that works well in total darkness. Digital cameras automatically correct electronically and mechanically for varying ambient light conditions. Most can correct for background lighting conditions that would otherwise cause subjects to be underexposed, for sunlight glaring directly into the lens, and for unusual brightness that would otherwise wash out a useful image. All of the images can be viewed in real time, and one or more persons can simultaneously control the views and scroll back to freeze and enlarge frames. All viewers can adjust the brightness, contrast, colors, and zoom, and can apply filters to bring out foggy details. Any viewer can save anything viewed.

All but the smallest indoor cameras now include a microphone so that both video and audio are recorded. Outside the facilities, some installations interface with an emergency intercom so background sound is continuously monitored and dialog recorded. Emergency intercoms are particularly good deterrents because security personnel can confront possible troublemakers and advise them that they are being recorded. Most will back off or flee before trouble occurs. Although many camera control systems provide electronic zoom, the primary zoom function should be mechanical. Electronic zoom systems are inexpensive, but they lose resolution and images can quickly become unidentifiable.

Miniature cameras can be particularly useful and save money as well. Cameras less than one inch in diameter can provide sharp detail. Some units with zoom lenses and pan and tilt mounts can be hidden within any common object.

Opinions vary as to the ethics of using concealed cameras, and some state laws limit their use as invasions of privacy. Generally, the rules of good surveillance practice are the same as monitoring telephone conversations, email, and Internet use. In order to protect itself, the organization has a right to see and hear whatever is going on in and around its premises. But signs should be posted so that workers and the public are advised that all persons entering the premises may be monitored for safety and security purposes. Company policies should explain this, and why surveillance is necessary.

23.7.2 Camera Locations and Mounts. To serve as a deterrent, some cameras should be readily visible, but this may present some problems. Overt cameras placed too high are not good for identifying people, while cameras placed too low can be spray-painted, covered, or smashed. Any visible camera can be shot out (perhaps unnoticed using a silencer). Older cameras are good for overt use because they are large and may provide backup surveillance. Another option is dummy cameras, some of which have visible indicator lights or automatically pan to attract attention. Some dummy cameras are in fact alarms that will trigger if the unit is disturbed or if its wires are cut.

Surveillance of a wide area requires an elevated camera that can see greater distances, zoom in on details (especially faces and vehicle license plates), and closely follow an event. Wide-area views are also necessary to spot multiple troublemakers. Several camera angles will likely be needed to gather good evidence.

Outdoor cameras are usually attached to a swivel, tilt mount, and inside of weatherproof domes or larger enclosures when large zoom lenses are needed. There is usually no attempt at concealment because the cameras are too high to reach, and they provide good deterrence as well. Many may be inconspicuous, however, or even disguised. When large cameras and enclosures are used, the direction the camera is looking in may be obvious. Often, troublemakers will create a distraction to lure the cameras away from the real problems.

23.7.3 Recording Systems. Once VCRs were the norm, but now most recording is done with hard drives that can store hundreds of hours of audio and video taken from multiple sources. Increasingly, solid state drives are also used.

Once, only one person could review the tapes and then, usually, only after the recording was complete. Now one or more people may access and analyze the information simultaneously. Once, 31 VCR tapes were recommended for each system: one tape for each day of the month, after which the tape was inspected and reused. Now one hard drive can record a month's data, which can then be inexpensively archived on DVDs or at remote storage locations. The VCRs often recorded in a lapse-time mode to save tape until an event occurred to trigger real-time recording. Now everything can be recorded continuously in real time to examine fully the happenings before, during, and after an event.

23.7.4 Camera Control Systems. Camera control systems commonly can direct the swivel (pan), tilt, mechanical zoom, aperture, and background lighting of each camera. Some control systems can automatically pan cameras back and forth across large viewing areas and also provide motion detection. Some control systems

can also stop panning and zoom in on any unusual event. This is a valuable feature that must be carefully programmed so the system is not distracted by spurious events or deliberate diversions. Usually, each camera can also be controlled manually and often from remote locations.

A major advantage of a digital system is that each camera provides continuous images, usually at about 30 frames per second, which fiber optic and other broadband connections bring to the recording system in real time and high resolution. The control systems usually allow one or more persons to roll back the images without affecting the real-time recording. Each viewer can scroll backward and forward, freeze frames, zoom, crop, or enhance the image electronically, and save any material to another medium or system as evidence.

23.7.5 Broadband Connections. The advantage of broadband camera connections is that the information is available in real time and with high resolution, and that more camera control functions are possible. Fiber optic cabling does this best, but it should be dedicated and well protected. Remote cameras can also be connected via a LAN, wireless, or broadband Internet connection using Transmission Control Protocol/Internet Protocol (TCP/IP), if done carefully, so that the data are secure and other traffic is not impeded.

Digital multiplexing is another advantage of broadband connections. Multiplexing over metal wiring once necessitated delays, lower resolution, and fewer frames transmitted per second, even though each camera was providing continuous, high-quality images. Now digital multiplexing over fiber connections allows all data from all cameras to be transmitted simultaneously. Multiplexing is accomplished remotely as signals from several cameras are combined into a single broadband connection to the control system. Connections among control systems are also multiplexed to minimize line charges. The signal interfaces between each camera and the multiplexer can be fiber optic, microwave, a network, or the Internet. This way, hundreds of cameras can be networked economically.

Another benefit of today's TCP/IP connections for cameras is that they can be Internet-enabled, allowing monitoring stations to receive data without having to retrofit special-purpose cables throughout the structure. A disadvantage of such systems is that they are therefore susceptible to the same attack methods that are developed to hack any other Internet-based systems. Proper use of encryption (e.g., virtual private networks) for camera output and controls can help protect the surveillance systems against electronic tampering. Encryption can also interfere with the *playback attacks* so popular in movies, where criminals insert devices to cause an infinite loop of playback and fool observers into believing that all is well in the compromised surveillance zones.

For additional information on penetration techniques, see Chapter 15 in this *Handbook*.

Radio or microwave multiplexing is another option, but it is not very fast and is subject to interference or disruption during severe weather conditions.

23.8 FACILITIES DESIGN. Good protection begins with good facility design that will ensure the safety of the information infrastructure and the people who use it. Protective systems become expensive and inefficient when the premises themselves do not facilitate good security. Proper interior design, facilities layout, engineering, and construction can maximize security effectiveness, minimize costs, enhance productivity

and morale, and generally boost profits. The starting point is an inspection of all sites and review of all as-built plans and construction documents.

Premises security inspection and review are best augmented by using independent, outside security experts. Comprehensive and objective architectural, design, engineering, IS infrastructure, and premises security experience are all needed and may not all be available internally or from vendors. The inspection and review process always must be threat-specific and must relate to a thorough threat assessment (which is described in Chapter 22). The premises inspection can then serve to validate the vulnerabilities identified by the threat assessment. Once some of the vulnerabilities are corrected, the premises inspection should be repeated.

Effective infrastructure protection can prevent trouble from happening. To do this, there must be ample early warning, which, in turn, requires good facilities planning and design, effective premises and infrastructure security in place, and the awareness and vigilance of everyone in the area. All must work together seamlessly, efficiently, and proactively.

Good facilities design can be efficient, nonintrusive, cost effective, and inexpensive—even within existing facilities. Here are some guidelines and suggestions.

23.8.1 Choosing Safe Sites. If possible, the building site for any facility that will house IS infrastructure—especially for highly secure applications—should be far from obvious manmade threats such as fuel depots, railway lines, suspension bridges, high-speed highway overpasses, and airport flyways. Sites should avoid obvious natural threats such as floodplains, areas subject to forest fires, active volcanoes, and old underground mines subject to subsidence of the overlying areas.

Access to the site should include provision for emergency evacuation of employees and visitors and access for emergency vehicles to reach the facility even if the main public access is blocked.

In cities, the building should not be situated in high-crime neighborhoods or in industrial parks with inadequate security supervision or histories of theft, vandalism, or sabotage. Decisions should include attention to public records indicating sufficiently competent police and fire-response teams; interviews with other users of industrial or commercial buildings in the area may be helpful in making decisions about the final choice of location.

Sites for equipment rooms and utility closets should be protected from possible threats. Infrastructure sites should be located far away from all piping, tanks, or equipment that uses any liquid that could possibly leak or spill. Most plumbing tends to leak unexpectedly at some time or can be intentionally breached, so it is well to assume that any pipe, connection, container, or pump will eventually burst or leak. Placing sensitive sites at a distance from such threats is safer and cheaper than any other form of protection. The danger zone where leaks may spread must include an ample vertical distance and horizontal area. Begin the danger zone with a pyramid from any leak source, downward several floors, and outward horizontally well beyond the infrastructure areas.

Most buildings require fire-suppression sprinklers in all areas. Any may activate by accident, vandalism, or an actual fire, and may cause considerable flooding. Also, infrastructure sites should be located away from windows, exterior walls, roofs, skylights, building setbacks, and storm drains, which are all potential sources of flooding. Treated water used in heating, air conditioning, and solar systems presents a worse problem, in that the chemicals can quickly destroy electronic equipment, connectors, and wiring.

If all of the infrastructure cannot be positioned at a safe distance from all liquids and hazardous materials, there must be special protections and alarms. Protections include sealed, waterproof construction, drains, berms or other diversion devices, and proper materials close at hand to control any spill. There should be environmental alarm sensors near where leaks or spills could occur. Floor drains must also protect the equipment areas, especially those that use sprinklers; otherwise, cleanup will be difficult.

Infrastructure sites should not be visible or accessible from any public area. Infrastructure wiring or cables should not run through the building core, public, or mechanical areas, including rest rooms, slop sinks, janitorial closets, and stairwells. Avoid placing equipment or cables where any persons might loiter. All equipment room entrances should be clearly visible from workplaces where employees can readily observe all comings and goings. Choosing inherently safe sites and entrances greatly reduces both risk and costs because less security is needed.

For effective security control, there should only be one access point to each critical area, used for both entry and exiting. However, if local fire or building codes require a secondary means of egress, a delayed-access panic bar is usually acceptable. Such a system delays releasing the exit lock for a few seconds, while an alarm sounds and a surveillance camera is triggered. There should also be surveillance cameras with motion detection throughout all secured areas. All locked doors should look as alike as possible from the outside and be identified only by a room number that looks similar to that of any other premises door. No functional name, person's name or title, or any other means to identify what is inside a locked area should be apparent. No signage or directory should include a functional, personal, or departmental name, but only area or room numbers and directional arrows if needed. Only floor, room, or suite numbers should appear on premises signs, on floors.

23.8.2 Physical Access. Access to the facilities site should be controlled. Guard stations at the entrance won't stop terrorists, but they can provide early warning of trouble. All such stations should be under constant surveillance and have panic buttons installed for emergency notification of the security teams that there's a problem at the entrance to the site.

Reinforced concrete bollards can prevent physical attack on a building by interfering with even high-speed ramming by ordinary vehicles such as cars, vans, and light trucks. They will not necessarily prevent attacks using commandeered large vehicles such as loaded dump trucks.

The entrances to buildings should not include large glass walls, which could be demolished easily to gain access to entrance halls for detonation of bombs.

Physical access to all parts of the information infrastructure should be restricted. All information system and network equipment must be inside equipment rooms that are always locked and accessible only to those who need to be there. All utility and wiring closets, power distribution panels, patch panels, wiring blocks, and terminations should be located inside equipment rooms that are always locked. If possible, do not allow unrelated systems, equipment, or utilities inside a restricted area, so that a technician working on an unrelated system cannot access the information infrastructure. If this cannot be avoided, IS personnel should always escort others entering these areas. In high-security areas, all persons entering should be escorted, or use two-trusted-person teams where each person observes the other. Guards and facilities-security personnel are good premises-security escorts but probably do not know the infrastructure and are not the best choice here. IS personnel are better escorts, even though they may

not know all the security details. In any event, all persons entering must be positively identified and each visit logged.

It is wise to put critical electrical distribution panels inside an equipment room, so they are quickly accessible and also protected. This is a safety issue as well; someone working on equipment or wiring can readily see that the circuit breaker is off. Otherwise, electrical distribution panels should be located inside a locked area that is unmarked from the outside other than by a coded location number. Panels are often located in public areas in many buildings; they must then be securely locked and alarmed as well. Whenever an electrical panel controls anything critical, access to the area should be restricted and the room alarmed. These precautions reduce any loss of power to critical systems by accident or intentionally.

One of the cases used in training courses about facilities security is the unlocked patch panel discovered by a physical-security expert in the public hallway of a hospital; the director of the intensive care unit (ICU) was horrified to realize that the breakers controlling the life-support equipment for the patients in the ICU were open to any child or malicious person.

A mantrap can best control access to critical equipment rooms. A mantrap is a two-door arrangement with an inner door, outer door, and a short corridor in between. Both doors are interlocked; one must be closed before the other can open. The corridor usually is constructed with a shatterproof, full-height glass partition (often one-way glass) on one or both sides for surveillance. Both the doors are usually windowless. And each door must have a strong access control system. But for safety, the outer door can usually be opened from inside the corridor using an alarmed panic bar. Usually, one or more surveillance cameras are positioned to identify anyone entering or leaving the mantrap.

Emergency intercoms within the mantrap corridor and at the entrance and exit points are strongly suggested. Conspicuous duress/assistance buttons should activate silently, so security personnel can monitor the area and speak with or challenge anyone who cannot pass through properly. The alarms are best silent so that anyone under duress is not further threatened. Emergency intercoms often include small cameras that are inconspicuous or concealed, so that security personnel can see what is happening, assist if someone is ill or somehow becomes trapped, and avoid public-safety issues. A well-designed mantrap can be valuable for all but heavily traveled entrances.

Mantraps can tighten security in many ways and are usually not intrusive. They can detect *piggybacking*, when one or more extra persons closely follow someone who is authorized to enter. Mantraps in themselves do not preclude propped-open doors, but they make removal of objects from the room difficult and risky. The access control log and surveillance recordings can identify troublemakers and provide strong evidence to convict those who might otherwise be suspects or persons unknown.

23.8.3 Protective Construction. The design of a new building can improve physical security.

- Avoid vulnerable construction methods such as using structural columns in large, open spaces on the ground floor; destroying such columns using explosives could cause serious damage to the structural integrity of the building.
- Don't allow architects to include decorative indentations on the outside of the building; such vertical patterns can provide the means for rock-climbing attackers to climb the walls easily.
- Place all IS in a hardened, isolated area.

Equipment rooms require sturdy partitions for good security and to support the considerable weight of wiring and equipment. Moreover, the partitions and walls must remain safe and stable during any seismic activity, such as heavy road traffic or a sonic boom, explosion, or earth tremor. Floors may also need to be reinforced. Sturdy partitions deter forced entry, which may be otherwise accomplished with little more than a pocketknife and a fist. Existing walls, partitions, and floors should be inspected and any subsequent alterations approved by a structural engineer. Consider too what might be needed long into the future; changes later may be very costly or simply impractical.

By definition, windows are not sturdy partitions. Avoid allowing nonsecurity personnel to pressure architects into including large windows, allowing visibility and possible intrusion into secured areas. Allowing uncontrolled visibility into operations centers may allow malefactors to learn more than they should about personnel, procedures, and schedules simply through observation.

Security doors should be sturdy also. They should be metal, fire-rated, and relatively heavy and use heavy-duty hardware. Try not to call attention to controlled doors by any distinctive external appearance. If many occupied areas use wood doors, the security doors should not stand out. Use wood-faced metal doors and hardware that looks similar to all the other doors. Sometimes secure-looking dummy doors that lead to nothing important are used for deception or as an alarmed honeypot to draw troublemakers away from secure areas.

Well-constructed partitions and ceilings will also seal out smoke and contain smoke in the event of an interior fire. Weatherstripping around the perimeter of each door is recommended to keep out dust, contaminants, and humidity and to trap any smoke.

If possible, wiring that must run inside a door should be routed within the hinges so that no wires are visible at any time. Exposed wires can be damaged by accident by cutting, or compromised in many ways. The major hinge manufacturers can supply special hinges to conceal most wiring and match the appearance of other premises hardware.

Masonry partitions are usually unnecessary unless there are special structural or very high security requirements. Drywall partitions with metal studs are usually sufficient but should be extra sturdy. Type-X fire-rated drywall panels at least three-quarters of an inch thick are recommended. Better yet, use double half-inch- or five-eighths-inch-thick panels. Existing drywall partitions can easily be double-paneled for added security and strength. Masonry partitions, especially, and often drywall as well are usually faced with sturdy fire-rated plywood for attaching equipment supports and wiring. Whether the plywood is mounted on stand-offs so that wiring can be run behind it is a matter of preference. Usually, it is more efficient to run all wiring exposed inside a secured area.

Do not use a suspended (hung) ceiling in any equipment room. Suspended ceilings just add cost, inconvenience, and diminish the volume of the room. The plenum space above the suspended ceiling is a fire hazard, and everything inside it must be fire rated, including any exposed cables, which adds unnecessary costs. Most building codes require separate fire detection zones and suppression heads within every plenum, which is a major cost. Remove any suspended ceilings and get more useful space for cables, better air circulation, and easier maintenance.

Avoid raised floors for the same reasons. They are very costly, functionally unnecessary for most installations, and take up extra floor space for ramps. Like suspended ceilings, raised floors create a plenum space that needs separate fire detection zones and suppression heads. Raised floors usually restrict the ceiling height because the space was not designed for this, so the room must be enlarged to accommodate everything,

and often with little expansion provision for the future. Raised floors soon become dust traps and, in time, usually a clutter of new, old, and abandoned cables that no one can figure out. Many equipment failures have been caused by overheating due to airflow restricted by too shallow a raised floor or by too many obstructing cables. A raised-floor plenum is rarely needed to supply cooling air. Surface-mounted ducts and registers usually can do the job better and much cheaper, and are easily cleaned and modified. All of the wiring can usually be routed efficiently above and between the equipment or by using inexpensive cable troughs mounted on walls and ceilings. If needed, floor outlets can access trench ducts in the floor, which may already exist, unused. Conduit can be installed through the floor slab and along the ceiling below for special needs.

It is important for all wiring outside of equipment rooms to be protected inside of metal conduit. This conduit should not be thin-wall or plastic but rugged, heavy-duty metal. Thick-walled metal conduit is strong, harder to damage or breach, and provides good magnetic and electronic shielding as well. Metal conduits can easily be, and should be, well grounded. Obtain expert advice on where and how to connect the grounds to avoid interference. Metal conduit also may serve as a proximity sensor to warn when something gets too close. Alarm wire concealed within the conduit offers early warning of trouble and often can pinpoint where it is located. Sometimes an inert gas or gel is used to pressurize conduits to protect the wires; a pressure drop indicates a leak or a breach but does not indicate where it may be.

All conduits should look alike, whether they carry power or data or control security systems. Although the diameter of the conduit must vary, the general appearance should be the same. Do not label or mark the outside of any conduit except with an alphanumeric code. Cables inside of a conduit should not be marked either, except alphanumerically. Generally, any wires within conduit should emerge only inside of a secure closet, equipment room, or junction box, where the wires should be labeled with appropriate codes that can be deciphered using up-to-date electronic summaries available only with authorization.

Data cables and wiring are fragile, whether copper or fiber optic. Any undue pressure, bending, stretching, vibration, or crimping can alter transmissions or cause failure. Fiber optic cables are especially fragile; metal conduit is usually the best and least expensive protection. This avoids special sheathings that are often cumbersome and costly.

Any wiring, whether metal or optical fiber, can be improperly specified, installed, or terminated. Substandard wiring or installations may function well temporarily, but future failure is likely, possibly hard to locate quickly, and will be costly to fix. Therefore, it is important that all cables be acceptance-tested and certified by an independent expert before the installer is fully paid.

Critical cable runs should be equipped with alarm facilities from end to end, whether the conduit carries power, data, or security. There are several ways to alarm a cable run. Outdoor conduits are often pressurized with nitrogen to keep out humidity, or with a special gel to keep out oxygen and humidity and to stabilize the wires inside. Interior conduits can be pressurized and alarmed in the same way. Monitoring the pressure provides an alarm when trouble starts (including failed seals that must be fixed quickly). The system is effective and provides early warning, but breaches cannot be pinpointed, and any future wiring changes may be difficult.

Proximity and pressure-sensitive sensors also can alarm the entire length of critical cables. A monitored run of conduit can be very long and may continue through areas that are difficult to protect or offer no concealment. While surveillance and intrusion detectors can protect most vital areas, there is often much infrastructure that can be

protected only by sensor wires running the full length of the conduit. Mechanical pressure sensors will detect unusual vibration, while proximity sensors indicate a change in the magnetic or electronic environment surrounding the cable or conduit. Newer systems utilize fiber-optic sensors that monitor sound pressure. Some of these systems are smart enough to distinguish routine, harmless events from possible trouble, and many can roughly pinpoint the location as well. Sometimes the conduit itself is the sensor, or an external wire is attached to the conduit, but these approaches are often ineffective.

23.8.4 Using Existing Premises Alarms. Various codes require workplaces to have visible and audible fire alarms. And most workplaces have voice paging, emergency intercom, surveillance, and premises security systems as well. All of this equipment can be utilized effectively to augment and support information infrastructure security. Audible alarms are used when persons at the scene must take immediate action, such as to lock down or to evacuate. Conversely, silent alarms provide early warning and allow security to monitor the scene discreetly, gather evidence, and respond and assist as needed. All these alarms can be integrated into infrastructure protection systems to provide better early warning and extra layers of protection.

All alarms and alerts should be transmitted to a central security desk or guard station. The purpose is to document and manage incident response, summon assistance quickly, monitor the scene, accumulate evidence, and support all of the response resources. Central management is especially necessary when threats cascade or multiple events occur, as they often do. Security managers, IS managers, the infrastructure security manager, and key department heads also should be notified immediately. Some of these people may be offsite or at remote locations but will need to communicate effectively with at least the operations center. Notifications and the subsequent communications should be quick, private, secure, and logged to document the events. One or more online backup security control and operations centers provide redundancy, support and assistance, and strong security.

An effective method of premises-wide alert uses voice codes broadcast over a paging system. These are usually scripted and often prerecorded so that alerts can be initiated automatically, remotely or manually. Hospitals do this effectively with their color-named codes, equivalent to silent alerts, which do not seem unusual to the public. An effective system of alert codes in a large organization also uses the names of fictitious persons. In a smaller setting, such as a school, where the names of all personnel are known, there can be alert messages to an individual to take an innocuous action, which is understood to be an alert code. Additional alphanumeric information in a message can identify the general location of an incident. As in a hospital, it is well to add similar codes for other routine purposes, so the public will generally tune out all the paging. This system is particularly useful when violence is threatened or has erupted.

Although most security personnel now have portable radios, there may be many areas of no reception, and few use earpieces so others cannot readily hear what is happening. However, all security personnel must know immediately when trouble looms, and they must be alerted in a way so as not to excite others. Also, everyone inside the premises needs to know when an emergency threatens. Indeed, everyone has a legal right to know and to promptly receive instructions for their own safety. Anything short of this will result in considerable liability. Therefore, effective procedures, clear simple instructions, good preparation, and periodic training can protect everyone and provide strong security.

The widespread availability of mobile phones has added another channel for emergency notification. Employees can be notified of emergencies using the Short Message Service (SMS) for text messages. Using computer-generated messages or even manual messages, security personnel can quickly and easily reach many of the people involved in emergency response. If phones are configured to use a characteristic tone for such messages, and if the phones are not generally used for casual messaging, such arrangements may be useful adjuncts to other notification methods.

23.8.5 Clean Electrical Power. Protecting electronic equipment requires a source of electrical power that is consistently *clean*. Clean power usually means alternating current with tightly limited variations of amplitude and waveform. Power outages and brownouts (low-voltage states) can cause obvious trouble, but numerous other disturbances can disrupt or damage the information infrastructure. Some of these include dips and sags, spikes, transients, and magnetic or radio-frequency interference. Most of these are intermittent and not necessarily present when the circuit is tested. Understanding each term is not as important as knowing that a wide variety of problems commonly occur in power lines, randomly and without warning.

Brownouts are particularly harmful. These are voltage reductions by the electric utility that can cause air-conditioning equipment, cooling water pumps, and ventilating systems to malfunction or to shut down. The associated equipment may be damaged unless it is quickly shut down or switched over to an uninterruptible power supply or a motor generator.

Few power disturbances will destroy circuits or crash systems immediately, but most can cause cumulative damage. Each incident can weaken electronic circuits that will eventually fail for no apparent reason. Poor equipment quality is often blamed, because the cumulative effect was not recognized. Worse, replacement equipment will probably soon fail also.

Power disturbances can be measured to determine whether a particular circuit seems to be clean. Usually, a recording device is left in place continually for at least a week to measure and log the details of every event that occurs. (Avoid any test instrument that merely logs an unnamed event but provides no details.) An independent engineer who is not a vendor may best provide testing that is objective, comprehensive, and covers the entire facility.

Some circuits are likely to show intermittent disturbances caused by something nearby or within the building and sometimes by faulty wiring. Knowing what power problems arrive via the main service helps to determine whether the utility is at fault and to isolate where in the building other disturbances may originate. There may indeed be numerous causes of power disturbances, and all of them may be intermittent, which is why continuous seven-day, 24-hour monitoring is the minimum recommended.

Another major electrical problem is improper grounding that can damage sensitive equipment and cause interference in cables. Electricians must comply with national and local electrical codes, but they do not necessarily understand or provide the special grounding necessary for sensitive IS platforms. Most heavy equipment manufacturers require that each unit have an isolated ground connection with a dedicated wire all the way back to the central building ground bus. A few manufacturers do not provide installation specifications unless asked, and some installers disregard them to remain price competitive.

Opinions vary as to the best building ground configuration for information systems, and as local conditions can vary significantly, no one approach is best. It is wise to consult an independent engineer to inspect the grounding configurations, and to

recommend and certify local code compliance. It is also important to provide separate circuits for all IS equipment, and where the equipment plugs into a receptacle, there be no more than a single receptacle. Separate circuit breakers are usually required for all equipment that can draw high current, especially if the load may cycle on and off. This is required by code for large motors, such as pumps, air conditioning, and elevators, whose cycling can cause dirty power on other branch circuits. But copiers and large laser printers (especially older ones) can also create electrical disturbances when starting and when the fuser-heaters cycle. All types of lights can cause a dip or a surge when many fixtures are switched on or off at once. (The newer fluorescent light fixtures with electronic ballasts conserve power and cause much less interference.) A separate circuit connection somewhat isolates the hot and neutral wires from other circuits, but interference may be generated through the grounding connections that often are daisy-chained with many other circuits to cut costs.

Do not share a dedicated circuit with any of this equipment, which can readily disrupt and damage electronic equipment: time stamps, electric staplers, coffeepots, refrigerators, heaters, fans, or any other device with a motor or solenoid. Even if inaccessible, a dedicated outlet should be a single receptacle, not a duplex. It is all too easy for vacuum cleaners, power tools, or maintenance equipment to use the other half of the duplex outlet. This will cause a severe disturbance and may trip the circuit breaker. There must be plenty of convenience outlets that are readily accessible for all noncritical needs.

Yet another cause of power problems is excessive solar activity. These events can be measured only when they occur, which is randomly during an unpredictable interval of several years that peaks about every 11 years. Solar disruptions occur only during daylight hours. High solar activity occurred in 1988, causing major power outages and radio interference in Montreal, Canada. Daily solar activity reports, forecasts, pictures, and historical records are available at www.dxlc.com/solar and other sites. (See also Section 22.6.7.)

There are several remedial options when power disturbances are suspected, encountered, or even possible.

1. Eliminate the problem at the power distribution panel. Better grounding, more separate circuits, suppressors, filters, and isolation transformers may help. But this type of remediation can be difficult, costly, or unreliable.
2. Use a surge suppressor near the equipment it protects. This is inexpensive but useless against brownouts, outages, and severe disturbances.
3. Employ an uninterruptible power supply (UPS), which provides battery backup power, for each piece of critical infrastructure equipment. One type of UPS is activated only when the incoming power fails, although its battery is always being charged. A better type is always online, acting to filter the incoming power, suppress most surges, compensate for minor brownouts, and maintain full system power for five to 10 minutes following an outage—enough time to allow an orderly system shutdown. A third and best type of UPS always powers its load from its batteries, thus isolating the load from the power line and providing optimum protection.
4. For systems that draw high wattage, a motor-generator (MG) set will eliminate most power problems. An MG set is an electric motor driven by the utility power. The motor is coupled to a generator that will always supply clean power to the load. The generator is usually voltage-regulated automatically, although the

frequency can vary, which may disrupt some timing circuits. Usually, if there is a power outage, mechanical momentum of the unit will provide sufficient power for an orderly equipment shutdown. Motor generators are still used for the ultimate in filtration and regulation, but there must be an electrical bypass to facilitate maintenance.

UPS units should power and protect the servers, network and telephone equipment, computers, and critical monitors. Most UPS units provide outlets with no backup power but with noise suppression for printers, transformer "bricks," fax machines, and peripherals that do not need to be kept running. Most of these devices are somewhat expendable and quickly replaced by spare units if one is damaged.

Individual UPS units, placed near the equipment they protect, cost less and can be powered off by the operator to better protect equipment that may be vulnerable, even when the equipment is already shut down. This can provide an extra layer of protection where lightning might strike. Larger UPS units are used in equipment rooms where they can also monitor and log all power events and trigger remote alarm indications. Most good UPS units can initiate an automatic equipment shutdown when the power fails, the UPS batteries are low, or someone intervenes manually. There can be an issue, however, when some of the protected equipment cannot be restarted until the UPS batteries are fully recharged. Some UPS units are also network devices that can report their condition to a remote location.

Many UPS units also provide telephone and Ethernet line filtering and suppression, and these features should be used if possible. Lightning transmitted over communications wires can readily damage telephone instruments and modems. Power and communication line spikes can occur asymmetrically, and can devastate equipment that one disturbance alone would not damage. A good UPS unit with communications line suppression is best able to stop both types of spikes.

Another benefit of UPS units is that when an emergency generator is on line, the electrical power is usually much dirtier than normal. The extra filtration and stabilization provided by the UPS units may be the difference between having IS equipment operating or crashing.

23.8.6 Emergency Power. Most of the critical systems and infrastructure must remain fully operational during any electrical power problem. Filtration and suppressors cannot compensate for power outages, handle most brownouts, or cope with major electrical interference. Disturbances might come from lightning (even when it is too distant to be seen or heard), severe solar activity, or radio-frequency interference corrupting the utility power. UPS units can deal with some of these conditions, but only briefly. Therefore, a backup emergency generator may be the only way to continue operations during sustained power problems.

Although backup generators are often the only alternative, they are not a panacea. Generators are expensive and complex to install, require at least monthly exercise, and are not always reliable. Their voltage regulation is marginal; during sudden load changes, the output voltage and frequency may fluctuate as well. As the load increases, more current is drawn, and, if the generator is overloaded, the voltage will drop and the frequency may drop below 60 Hertz (which can disrupt IS timing circuits). Because the load current increases to meet the power demand, the amount of heat generated by the equipment being powered will increase as the square of the current. Much of the IS equipment powered is inductive, and there can be a large starting power surge when it is turned on or restarted. Generators, therefore, must have ample reserve capacity

beyond the anticipated equipment loads. And given their cost, generators should have ample reserve capacity for the future as well.

Another issue is whether a particular generator can provide sufficiently "clean" electrical power to operate IS equipment as well as power for the other emergency needs of the facility. Be sure, therefore, that the generator specified has ample capacity and that it is intended for use with electronic equipment. Even then, interference from large motors or lighting systems can affect the electronic equipment.

Because backup generators are expensive and complex, planning is often shortsighted, and many installations are not well designed or adequately tested. The inevitable result is that many generators do not perform as expected. Here are a few examples of what can be overlooked.

After considerable discussion at a major money-center bank in Manhattan about what seemed to be the excessive cost of a backup generator, the project was begrudgingly approved. The generator was to power two identical computer systems running in tandem to support a critical securities trading operation. Because the generator would actually cost more than the two computers, cost was an issue until the bank realized that the generator would pay for itself within one day of use. Soon after completion, a sudden citywide blackout erupted and the outage lasted for three days. Despite much inconvenience carrying fuel from the elevator up a flight of stairs to the rooftop generator, the unit performed flawlessly—one of the few systems in Manhattan that did.

Many other generators did not start or cut over properly, despite warm, clear weather conditions. And others did not support the necessary infrastructure. Some installations did not think to include the power requirements of HVAC, so the computers had to be shut down within a few minutes to avoid overheating, even though there was ample electrical power for them. Generator power for other necessary support functions was neglected. These included network components and communications systems, lighting for users, an elevator for access by key people and to carry fuel to the generator, security and access control systems, and at least basic electrical provisions for food and rest for those keeping the vital systems running. Very few businesses thought to include all of the necessary support functions on their emergency power systems. This incident happened some years ago when power outages were considered very unlikely in Manhattan. Generators then were somewhat of a status symbol. But today, sudden blackouts anywhere are far more common.

A related example of shortsightedness occurred in a large suburban luxury hotel operated by a prestigious hotel chain. Following a severe thunderstorm, the power utility advised the hotel that they must lose power for several hours to repair a damaged substation. Given ample notice, the hotel set out hundreds of candles in the corridors, dining, lounge, reception, and pool areas, and started their emergency generator, which then cut over automatically as soon as the blackout occurred. Emergency exit signs and emergency lights in the corridors and stairs all worked properly. As expected, their batteries soon died but the candles functioned long and well. The generator also powered one elevator, the computer and telephone systems, the access control system, and all the room locks. The generator performed as expected, but the emergency response process did not.

Even with ample warning, no one thought to shut down the other elevators or to post signs to use the stairs. Two very frightened people were trapped in the dark, between floors, proving that a generator can be a liability and not a benefit unless operating procedures are carefully planned, well implemented, and periodically reviewed. There should have been a security checklist used whenever the generator started.

Another recent example involved a state's emergency operations command center, designed to remain safe and fully operational no matter what events might occur. A large generator powered all the critical systems. Everything had been tested many times and had operated smoothly as expected. But then trouble came during a heavy thunderstorm in the vicinity. Electrical power for most of the city flickered several times and then returned to normal. However, the generator tried to start at the first sign of trouble and then faltered as the power returned. A few seconds later when the power again flickered, the generator system had been damaged and was unable to start. The state was lucky that the generator was not needed then, but it was out for several days for repairs.

Most power failures begin with flickering and momentary outages, which can incapacitate a generator system that is not set up properly. Most mission-critical generators are set up to start the engine automatically, and many transfer power automatically as soon as the generator comes up to speed. Manual start-up and transfer are more reliable and cheaper, if trained personnel are always available. The best way to sequence automatic operation follows.

1. After the first start-up signal, the start-up sequence must continue until the engine starts, a failed-start timeout occurs, or the sequence is terminated manually.
2. Power does not transfer until the generator is fully up to speed, at a reliable operating temperature, and the utility power is unusable. All three conditions should occur before transfer, and there can be manual overrides as well.
3. All transfers back to utility power and the generator shutdown should be done manually. It is best also to be able to transfer each circuit individually to utility power.

There are countless examples of critical backup generators failing to operate as expected. Here are some suggestions to determine whether a generator is necessary for protecting information systems and how to utilize a generator efficiently and economically.

- Investigate the outage history of the utility feeders that serve the premises. The electric utility can usually provide this data; if not, the state's Public Utilities Commission usually can. Be sure to ask how the terms are defined, because an "outage" may only include interruptions that continue for more than several minutes. Also ask whether more reliable feeders are available. Loop feeders that are powered from both ends are more reliable and often serve critical equipment. Ask whether the distribution transformers isolate and filter out power disturbances, whether they can also regulate the incoming voltage, and, if so, the specifications.
- Find out which other customers share the same feeders, and visit them to discuss their experiences and to determine if they use heavy machinery. Although some safeguards are possible and may be at little or no cost to the utility customer, past history is not always a reliable guide to the future. The distribution grid changes as more heavy loads are added. Today, the threat of extended power problems is far greater than in the past and is increasing rapidly. UPS units, motor-generator sets, and backup generators may all be a necessity in mission-critical applications.
- Determine which of the IS infrastructure components need backup power from an emergency generator. Most critical information systems, equipment, networks, and infrastructure must be at peak performance at all times. And so must all the

office areas, support systems, utilities, and personnel needed to operate them. Outages can drag on for days or weeks with key people isolated and living inside the facility to keep the systems running. The generator power must serve all of these needs.

- Consider these support systems that may require emergency power:
  - All the IS security systems, protection and monitoring devices; perimeter surveillance, and access control systems, the security stations and consoles.
  - Fire stairs (which may become the primary means of entry and egress), emergency exit doors, fire alarms, and intercoms whose batteries will quickly discharge. Also the need for these batteries to begin recharging immediately as soon as backup power is available.
  - Heating, ventilation, air-conditioning (HVAC), and process-cooling systems, including all the controls, fans, pumps, and valves needed to operate the critical and support systems. In addition to equipment cooling, it is best to provide room comfort for users, operators, and administrators. Area air conditioning may not be possible, but at least supplementary heating in winter and adequate ventilation will be needed.
  - Sufficient lighting for key personnel, equipment rooms, utility closets, corridors, rest rooms, and food service. Many individual light switches can conserve power and generator fuel. Battery-powered lights are suitable only for immediate emergency egress and cannot provide area lighting.
  - Enough live convenience outlets for test equipment, work lights, and any accessories that must be used. Live receptacles may also be needed for portable fans.
  - Sufficient food service equipment and refrigeration, running water and sanitary facilities, and a sleeping area for 24/7 operations that may have to continue for several days.
  - An elevator for critical access to the site, for medical emergencies, delivery of food and supplies, and to carry fuel for the generator.
- Compile a list of all the items a generator must power. Then total the rated power of each item to determine the size of the generator and the number of circuits needed. Power ratings for equipment usually are shown on a nameplate near the power connection and listed in the instructions. Ratings may be given in watts, amperes, or volt-amps. Generally, watts and volt-amps are assumed as equivalent. The latter value is the product of multiplying the rated voltage (e.g., 120 volts) by the rated amperes, while the former multiplies that number by the equipment's power factor. Large generators are rated in kilowatts (1,000 watts) of power. Units intended for short duty cycles cost less and may fail during prolonged, continuous duty. An experienced engineer should review this process.
- Consider the costs per hour of lost productivity, business, and goodwill if any information systems are disrupted. Add to this the recovery costs. In the example of the bank given earlier, the first day that the generator was needed saved the entire cost of the backup power system. The second and third days of that particular outage were sheer delight to the bank as most of their competitors faltered.

Electric codes may require, and good practice dictates, that a generator be sized to handle the sum of the individual start-up loads. This may seem wasteful, because not

all loads start up at once and an average operating load will be somewhat less than the sum of the parts. It is nonetheless a wise practice to provide for the maximum rated load, with additional spare capacity for improved reliability and future additions. There are several reasons for oversizing the generator. When power is first transferred to the generator, the sum of the initial surges can far exceed the anticipated load. All of the UPS units and other battery-operated devices will begin to recharge, and all equipment motors may concurrently draw their maximum surge currents. Extra generator capacity ensures a smoother transfer with better voltage regulation and enhances the system's reliability.

Most large generators produce three-phase power. And each of the three outputs should be balanced so that each "leg" draws about the same power. To do this, heavy motors, multiple light fixtures controlled by one switch, and other surge-producing equipment may have to be divided among the three legs. Existing wiring at the distribution panels probably will need changing to balance the legs. It is desirable, but not always possible, to reserve one leg for clean power to critical single-phase electronic systems. Balancing each leg is a tricky business best done in consultation with an independent engineer.

As many electronic systems as possible that are powered by a generator should also be protected by UPS units as discussed earlier, even though these add to the generator's load. There will be large voltage surges, dips, sags, and over-voltage conditions as the generator loads are switched and constantly change. Power disturbances will be much greater because electrically noisy motors and lighting cannot be isolated. The UPS units should include noise suppression and voltage regulation as well. And even with all of this, IS equipment will be stressed and may fail.

Locating the generator is the next challenge. The choices are on a roof or setback, inside the premises, or outdoors. Each site has advantages and obstacles. Outdoor generators can be the easiest and cheapest to install but also more expensive to operate. Outdoor generators are noisy, often unsightly, subject to vandalism, and local ordinances may restrict them. When located outside, weatherproof housings are needed to protect the engine, generator, and fuel tank. Most engines used outdoors need to be kept heated, which can become a high overhead expense. Noise is another problem, and persons nearby may object. It is important to use good mufflers and to get written permission from nearby property owners and other tenants. Outdoor units should be fenced with plenty of room for maintenance and fueling. A generator shed is best, if possible, but this does not reduce the need for heating and a good muffler. The whole installation should be securely locked and protected by an open-door alarm and motion detectors, and be in the view of surveillance cameras. Floodlights may deter vandalism and will assist refueling.

Generators on roofs or building setbacks present other problems, and these installations too may be restricted by local codes. The first problem is weight. Structural reinforcement probably will be needed. The next problem is getting the unit in place, which may require a crane or a licensed rigger. Very few building elevators come up to the roof level, and they may not be able to handle even a disassembled generator's parts. All the generator components may have to be rigged up outside of the building or manhandled up the fire stairs.

Installations on top of building setbacks will need a special access door, and moving heavy equipment across a finished floor requires heavy planks and floor protection (e.g., sheets of Masonite or plywood) under the casters to avoid considerable floor damage. There must be sufficient space on the roof or setback to fuel and service the generator safely. Noise will usually be a problem and vibration as well.

Indoor installations offer both advantages and challenges. An indoor location that is sometimes feasible is a heated garage-type ground-floor room with a standard garage door to open when the generator operates. This arrangement is good because it is inconspicuous, fireproof, easily protected, and convenient for fueling and maintenance. And, should a generator fail, a trailer-mounted unit can be hooked up easily.

Inside generators may be prohibited by building or fire codes. Large rooms are needed to facilitate fueling and maintenance, and large ventilation systems to dissipate the considerable engine heat. The engine exhaust can be well muffled and piped outside, while engine-intake air is ducted in from outside. Heating and ventilating the room must be designed correctly for both very hot and very cold weather. The room must be fireproof and soundproof with fire alarms and a suppression system that uses chemicals or dry-head sprinklers that cannot freeze. The floor may need reinforcement and vibration isolators. A floor drain is advisable and must be environmentally approved.

There are advantages to indoor installations. The generator and its fuel can be kept warm easily. Starting is easier and more reliable. Fueling is easier without having to brave the elements. There is less chance of water in the fuel, which can be fatal to diesel engines and maintenance is much easier.

Problems with building installations include building codes that allow only small day tanks for fuel. Every few hours, a lot of fuel must be carried in to keep the generator running. Fuel cannot be stored inside most buildings, and an elevator may not be running or available to help bring in fuel cans.

There are many possible fuels for emergency generators. Diesel fuel is the most efficient, and diesel engines can operate continuously for days but are hard to start, especially in cold weather, and cannot be hand cranked. Home heating oil is basically the same as diesel fuel and can be substituted at any time that diesel fuel is not available, but this requires extra fuel filtering.

If liquid fuel is used, the fuel tank should be full at all times to avoid condensation. Fuel additives can prevent gumming and assist starting. Make sure all diesel fuel is treated for use in a cold climate. Refiners normally do not use this process except in winter, but untreated diesel fuel turns to a gel near freezing temperatures and the fuel will not flow. Never let a dealer "cut" diesel fuel with kerosene, which is corrosive. Diesel fuel also requires additives to avoid bacteria buildup that will clog fuel lines. There should be OSHA-approved cleanup materials ready for any future spills or leaks.

Natural gas or propane are the most convenient fuels. Either one eliminates the day tank and refueling. These engines are the least polluting, and they start much easier, require no preheating, and can be hand cranked. But most are not designed for prolonged continuous duty. Gasoline engines are prohibited by many building codes and are rarely used except for small, portable generators. Gasoline is far more dangerous to handle and store, and gasoline engines do not hold up well under heavy loads.

Continually monitor the engine oil level and be ready to add oil as soon as it is needed. Most generators automatically shut down when the oil level is low. Some also shut down when overheated. Any unexpected generator shutdown will be catastrophic, so monitor closely for early warning signs of trouble.

Once the desired size and type of generator is decided, there are other considerations:

- Automatic engine controls and load transfer switches can be unreliable and may cause damage. Avoid these if possible. However, generators can be monitored and controlled remotely, as well as on site.

- Automatic starting can be unreliable. If the engine does not start quickly, the battery will quickly discharge, especially diesel engines in cold weather, which require glo-plug heaters. If at all possible, someone should be present during the starting process, using a checklist to verify proper operations and then transferring the load manually when the generator is ready. Switches that automatically transfer the load are expensive and sometimes fail. Always transfer back to utility power manually, and do this only after sensitive systems are put into a standby mode. Automatic transfer can cause major damage if the utility power flickers and goes out again or if the voltage or the frequency fluctuates during transfer, as it often does. Do not shut off the engine automatically. This is best done manually, and not until utility power is flowing smoothly.
- The best transfer switches allow each of the major circuits to be transferred individually to minimize the inevitable fluctuations likely to occur when everything is switched over simultaneously.
- An emergency generator must be exercised regularly. The manufacturer will specify when and how the units should be exercised. Usually, this must be done monthly and at medium to heavy load. When critical systems are involved, good security practice is to exercise the generator weekly. There should be a written, initialed log entry for each event, including each exercise, inspection, maintenance, oil check, and refueling. Always log operating hours.

Despite the cost and complexity, there is a great feeling of contentment in having a good emergency generator system that functions smoothly, especially when other organizations may be floundering. Once the generator performs well during a real emergency, even skeptics realize the value added.

23.8.7 Environmental Control. Even though today's information systems do not need as much cooling or the precise environmental controls that legacy systems once demanded, good control of temperature and humidity, good ventilation, and clean air are still important. Information systems can function reliably only when proper environments are maintained in both equipment rooms and user workplaces. But each area requires a different approach.

Air conditioning is basically intended to cool people; equipment should be cooled by a functionally different system, which is best called process cooling. The systems should not be intermixed, nor should either one substitute for the other. Building codes require HVAC within all occupied spaces, where people may congregate, or where there are workstations. Building codes also set minimum ventilation requirements for occupied space, including a minimum percentage of makeup (outside) air to be constantly brought into each occupied space so the inside air does not become stale. Most codes do not consider the needs of electronic equipment.

Electronic equipment has many special needs, and many are incompatible with the people comforts required by the codes. Most electronic equipment operates continually, whereas air conditioning operates mostly during business hours. Air-conditioning cooling systems may be shut down for maintenance, during a power brownout, off hours, or in cool weather. By contrast, process cooling must operate continuously and every day, so parallel and redundant systems are often used. The same air should be well filtered and recirculated with no makeup air added to introduce dust or contaminants. This also reduces the cooling capacity needed, so process-cooling equipment can be of smaller capacity and cheaper to operate.

Electrical equipment and wiring also need good humidity control, which process-cooling systems are designed to provide. These systems are designed to be easier, and faster to clean and maintain. Often many components are redundant and hot-swappable. Increasingly, the cooling unit is on the floor or ceiling of the equipment room, so that few ducts, dampers, or registers are needed.

All IS processing, storage, and network equipment should be inside dedicated equipment rooms, which also should be designated as unoccupied spaces to avoid the code-imposed air-conditioning requirements. Avoid using terms such as "computer room" or "data center," which are usually construed to be occupied spaces.

Both the process cooling in equipment rooms and the air conditioning in work areas must provide humidity control. It is important that relative humidity be controlled between 40 and 60 percent at all times, regardless of the climate or season.

When the relative humidity falls below 40 percent, which can easily happen in cold weather, static electricity is generated as people move about. Static charges can quickly accumulate to become many thousand volts, and a spark will jump to any object a person touches that is differently charged. Even though such a spark may not be felt, several thousand volts can annihilate electronic circuits. For example, a static charge jumping from a fingertip to a keyboard can cause serious damage to storage media and circuits. Much of the damage may not be readily apparent. Actual failure may be delayed, so the cause is not identified. Grounded strips can be installed on workstations, and service personnel should wear grounded wrist straps, although these do not completely stop the problem. The only effective solution is always to keep the relatively humidity above 40 percent.

Relative humidity above 60 percent also causes problems that will eventually destroy equipment and wiring. Above 60 percent, condensation and mold will begin to damage some components. Above roughly 80 percent, galvanic action occurs and will eventually cause serious trouble. The process is often called silver migration because most electronic connections are silver-plated. The phenomenon is similar to electrolysis (electroplating), but here the two metals are immersed in high humidity rather than a liquid and there is no external current needed for galvanic action to occur. Molecules of one conductor begin to physically move toward and attach themselves to another less active metal. Even though both surfaces may be gold or silver or copper plated, it is likely that they differ slightly in composition. Therefore, galvanic action will occur whenever the humidity is too high. Connector pins and sockets can disintegrate, fuse together, or fail electrically due to pitting. Printed circuits can also fail. Although this galvanic action happens slowly, it accumulates and is irreversible. The failures are usually without warning, and almost always, poor quality is blamed, rather than high humidity.

The only protection is to control humidity in both equipment rooms and work areas. Process-cooling and air-conditioning systems commonly do this by several methods. Both systems dehumidify naturally when cooling and can use a reheat coil to warm output air if the humidity is too low. Also when the humidity is too low, water is added using a spray, atomizer, or a wet screen through which the supply air is pumped.

There are additional protections, which are wise to install and maintain. In a cold climate, all areas and workplaces with electronic equipment should have low-static floor surfaces. This can be low-static carpeting or floor tile made for this purpose. Do not rely on sprays to control static electricity; they soon dissipate. Be sure that equipment room walls and doors are well sealed so that humidity, dust, and contaminants cannot migrate. Be sure the walls are well sealed from slab to slab and that the slabs themselves are impervious to dust and humidity.

23.8.8 Smoke and Fire Protection. Smoke and fire must be prevented within any equipment room because otherwise, considerable damage and disruption will occur quickly. No matter how small the incident, the effects of either smoke or heat are cumulative. Systems will eventually fail and usually without warning. Obviously, smoking tobacco and other substances must be prohibited—but equipment rooms can be tempting for untrained staff because these rooms may seem to be a safe, cool place to sneak a smoke—and they may even be largely unattended.

Equipment room doors should never be propped open by cleaning, delivery, or maintenance personnel, or others working outside when the air conditioning is off.

The first level of prevention is to keep everything combustible outside of equipment rooms. Paper and supplies not in actual use should be stored outside, never within an equipment room. Although most documentation is now made available electronically, reference materials or documents for emergency response when systems are down may be kept within the room; these should be stored inside fire-resistant files or cabinets for secure and rapid access in an emergency. There should be no trash receptacles within an equipment room, and shredders should be outside, under strict control. There should be a clear and firm policy that nothing combustible can remain inside an equipment room, and frequent inspections should be held to verify compliance.

Equipment and suppliers storage contents rooms should be designed to protect and not to accommodate people. There should be no unessential furniture within equipment rooms, especially desks that can become cluttered and that are not rated as fire resistant. A metal table with one small drawer and one or two metal chairs with fire-resistant upholstery are usually sufficient.

Inadequate firestops are a major threat that is often overlooked. A firestop is usually a sleeve and a special material to prevent smoke or heat from penetrating an opening in a partition, floor, or ceiling. It also stops the spread of flame. Many firestops are needed throughout the premises, including the building core and the mechanical, utility, and equipment areas—even within a one-story building. Firestops are rigorously required by most codes, but compliance is often inadequate and the devices often are breached by subsequent alterations or wiring changes.

Partitions, building walls, floors, and ceilings must all be fire-rated in accordance with the national and local building codes and other regulations. Proper construction usually is specified by an architect or engineer, and compliance is inspected or certified as soon as the construction is complete. Inspection often occurs before all of the mechanical systems and wiring installations are finished. Subsequently, if any penetration or opening is made through a wall, floor, or ceiling, its fire rating is thereby invalidated and should be recertified. The walls for all equipment-rooms must be slab-to-slab construction with no access for air transfer below or above the walls.

The Underwriters Laboratory (UL) or similar recognized authority rates and approves commercial firestops before they can be sold legally. Each manufacturer then specifies the approved applications, installation, and maintenance procedures necessary for compliance. It is therefore wise to utilize specialized vendors with extensive training and experience installing and inspecting firestops and to have them conduct periodic premises inspections and certifications.

Inadequate firestops are particularly common in the core areas of older buildings or where tenants occupy multiple floors. While proper firestops may have been provided during construction, installation of piping, cables, conduit, and subsequent wiring changes often breach them. Wires often poke through large, gaping holes hidden by a hung (suspended) ceiling or behind equipment racks or wire troughs. Proper firestops must be installed and inspected whenever changes occur. Many installers do

not understand this, or cut corners hoping no one will notice, or assume others will take care of it.

Without proper firestops throughout, fire and smoke can, and probably will, spread surprisingly quickly. And so will dust. There can be substantial liabilities if any people are harmed or equipment is damaged because of improper firestops or inadequate fire-rated construction. The costs, time lost, and reputational damage will be huge. Periodic and thorough fire inspections by an independent and qualified expert will quickly discover building and firestop violations.

Smoke is far more dangerous than flame. And all smoke is toxic! It contains high levels of carbon monoxide, which is invisible, odorless, and quickly fatal. Smoke is the product of the combustion of many materials, and most of these are dangerous to breathe. Some are immediately fatal. Even a little smoke can do considerable harm to humans and much harm to electronic equipment. Smoke is deceptive; even when there does not seem to be very much smoke or heat, and visibility looks good, people within or passing through the area quickly become disabled and some may soon die.

The first priority is the safety of people. Get everyone away from any smoke immediately, and keep everyone away. Only trained responders with proper protective clothing, equipment, and self-contained breathing apparatus should enter any smoky area. Generally, respirators are not enough protection and may leak as well. There must be no heroics; crawling through smoke on the floor or breathing through a wet rag are desperate measures that should be attempted only when unavoidable to escape the area. Everyone should wait in a safe place until firefighters arrive and then follow their instructions.

The best way to prevent an equipment room fire is to keep anything combustible outside the room. Documents, manuals, and supplies should be stored outside the room in closed metal cabinets. Inside furniture should be limited to a metal table and a chair or two. All waste receptacles should be outside. Once combustible materials are eliminated, the only smoke that develops will be from electrical overheating. Electrical fires rarely occur in an equipment room, and those that do occur are likely to be very small, brief, and cease as soon as electrical power is removed. (Note that most computing components now operate on five volts or less, so that a short circuit is no more dangerous than, for example, a shorted flashlight, which presents no smoke hazard.) While sometimes noticeably acrid, there is usually little visible smoke. Therefore, sensitive fire and smoke detectors and an effective means of fire suppression are needed and required by most building and fire codes. Good detectors can provide enough early warning to ward off trouble and injury.

Enough smoke or heat to cause actual equipment damage requires an electrical current higher than most components can draw. Circuit breakers and fuses usually will open before there is much smoke or damage. Perhaps the greatest risk is smoke from the ballasts in low-quality fluorescent light fixtures, which can put out considerable black smoke. Any smoke is corrosive and may condense on connectors and printed circuits, which may then eventually fail.

There should be smoke detectors in every equipment room that are connected to a central alarm system. There should be enough detectors to cover the entire volume of each room. Each detector should include an electric eye to look for haze or smoke, ionization sensors to detect products of combustion well before any are noticed by humans, and rapid-rise-in-temperature detectors in case there is enough heat buildup to cause damage. Even though detected, nothing will stop smoke generated by overheated wiring or components until the electrical power is cut off or other heat source is removed.

There must be a fire suppression system in every equipment room. Both code compliance and good security practice requires this. Fire suppression is best accomplished with sprinkler heads that spray water mist, even though some unprotected equipment may be damaged if the water is not effective quickly. Special waterproof protective covers are often kept near equipment in case of accidents such as ceiling leaks or a damaged sprinkler head. But if an area is already smoky, no one should attempt to place the covers.

Wiring, connections, and most components will dry themselves, even when soaked. The process may be hastened with lint-free towels and careful use of hair dryers. Keyboards, monitors, UPS units, power supplies, some disk or tape drives, and especially printers may be damaged and should be replaced until they can be inspected. Hard drives are usually hermetically sealed and unaffected. A few other components could be damaged by excessive heat, although water mist is very effective in quenching heat sources. Plenty of replacement items should be safely stored nearby. Handling damaged low-voltage components (such as most circuit boards) presents little risk to people—provided there is not too much water and the persons know what they are doing and how to avoid damaging the components. If in doubt, shut down the components temporarily.

Enclosed equipment cabinets offer the best protection regardless of the room's fire suppression system. Enclosed cabinets can monitor temperature and humidity, detect and contain smoke, sound alarms, and often contain systems to suppress a fire before trouble occurs.

Halon 1301 fire suppressant was once widely used in critical areas. But Halon is a fluorocarbon whose manufacture has been banned for many years. Today's chemical systems are designed differently; one example uses the FM200 Suppression Agent made by Siemens. The claimed advantage of the chemical suppressants is that humans can breathe the agent, at least while they are exiting the area. Another fire suppression system uses carbon dioxide, which is effective and less expensive, but can extinguish people as well as fires. The problem with all chemical agents, including carbon dioxide, is that they quickly mix with smoke and become very toxic. The agent itself may be safe to breathe, but the smoke mixed with it is not. These systems are also very expensive.

Regardless of the suppression system, there should be controls and a shutoff near the room's exit, but not accessible to a perpetrator. Generally, an audible, continuous alarm indicates that the suppression system is about to activate. There should be postpone buttons on the control panel, and perhaps remotely as well, that will delay activation for about two minutes while someone intervenes. The postpone mode generally pulses the audible alarm. A silent alarm indication should remain activated whenever a fire suppression system is disabled or the alarms are silenced.

The next level of protection utilizes several fire extinguishers. These are the most useful protections because the suppressant can be aimed where it is needed and not throughout the room. Carbon dioxide is best because it does not leave a residue. Chemical, powder, and foam extinguishers also work well but are hard to clean up. ABC-type extinguishers are best because they are effective for combustible materials, flammable liquids, and electrical fires, respectively. Several handheld extinguishers are better than a few large, heavy units. All fire extinguishers should be conspicuously wall-mounted or placed immediately inside and outside of entrances. An OSHA-approved red patch placed on the wall nearest to every extinguisher highlights its location. Also, check other OSHA, local code, and insurance requirements that may apply.

Supply air from the process-cooling equipment should be shut down quickly and automatically to avoid recirculating the smoke. The IS equipment may have to be shut

down soon thereafter before it overheats. It is best to shut down everything promptly and automatically in an orderly sequence—cooling, IS equipment, electrical power, and lighting—and then evacuate. Shut down the lighting, in case it is part of the problem. Shut-down should occur automatically with manual intervention from controls inside the room or remotely. Battery-powered exit and emergency room lighting are advisable so responders do not need flashlights.

A so-called crash cart is a good investment. This is used during a smoke condition, a water leak, and, it is hoped, before a fire suppression system activates. A crash cart is kept outside or nearby major equipment rooms and rolled to where it is needed. The cart usually contains covers to keep smoke and water out of racks and off equipment, large fire extinguishers, and sometimes respirators or self-contained breathing apparatus. The crash cart should include quick-reference procedures, and a checklist for protecting and shutting down the room, as well as safety and notification procedures—usually printed on plastic. The crash cart should be inspected and the procedures reviewed monthly, and there should be periodic training and exercises to practice using the equipment. Before the smoke and water covers are used, be sure the equipment is first powered off. Crash carts were important for yesterday's computer rooms but are increasingly unnecessary in a well-designed equipment room.

Finally, be sure to have smoke-exhaust systems available to quickly purge the areas of smoke. Most fire departments have portable purge fans with long fabric and wire hoses to reach outside. Do not allow anyone to use a respirator or breathing apparatus unless it is approved for this purpose and has been properly fitted to a trained person.

23.9 MITIGATING SPECIFIC THREATS.

Several other threats should be considered before good infrastructure protection is possible. Some of these situations are unlikely but potentially very costly if they should ever occur.

23.9.1 Preventing Wiretaps and Bugs.

Most wiretaps are placed at wiring junction points. Vulnerable spots are within equipment rooms, wiring closets, junction boxes, wiring blocks, or data receptacles. The tap wire can be fiber or coax or utilize a pair of unused conductors inside an existing cable. It is likely to be a small wire that is hardly noticeable, running to an inconspicuous place where monitoring and recording can occur. Once removed to a safe place, the data can be extracted by phone, wireless, Internet, or manually. Tapped data may even be encoded and stored on the victim's own network. Video and/or audio bugs used for spying are similar to wiretaps in that once the data is monitored, it must then be sent elsewhere for retrieval.

Unless all system data are encrypted—including all data, voice, and video traffic—wiretap protection must be strong because detection is difficult at best. First, determine which cables are critical and inspect the entire cable run. All cables should be inside of metal conduit. Data and power conduits should look similar and with no markings or labels except alphanumeric codes. Keep critical conduits as inconspicuous as possible, and away from places the public might access. There must be strong access controls, intrusion alarms, motion detectors, or surveillance where terminations, connectors, or wires can be accessed. Critical cable runs must be protected over their entire length. See Section 23.8.3 for ways to protect conduit and exposed cables.

Data cables between the desktop and wall or floor outlets are potential wiretap sites. Cables, harnesses, and connectors within office furniture systems may also be compromised. Reasonably good protection is possible with careful design, with devices that harden the data cabling against the possibility of a wiretap, and that detect disconnecting or tampering with any data wires.

For continued protection against wiretaps and bugs, even when all data are fully encrypted, there must be periodic and thorough visual inspections, sweeps for any unusual radio-frequency transmissions, and careful cable testing to determine any anomalies. Everything done must be logged and quickly analyzed. Unfortunately, most spying is never detected and can continue undetected at the will of the perpetrators.

23.9.2 Remote Spying Devices. There are very sensitive radio receivers that can monitor information system data through walls or from outside the building without the use of an inside bug or wiretap. These devices can simply listen to the data from afar. Such equipment is not available publicly and is well beyond the means of all but the best-financed spies. However, there are many such systems in use today, and many more will be available as prices drop. Any organization whose data are very valuable is a potential target. The best protection is good shielding around equipment rooms and thick-wall metal conduit for data cables, and everything must be properly grounded. There are also interference transmitters that may help; these broadcast white noise that can overwhelm signals radiated from the IS infrastructure.

23.9.3 Bombs, Threats, Violence, and Attacks. Violent events are unpredictable and potentially devastating. These are not accidents, but deliberate attacks, intended to disrupt, cause damage, and spread fear. The tragic attacks of 9/11 and their aftermath have proven the vulnerability of people and of their infrastructures. The vulnerabilities remain today, and the risks are even greater.

Protection against violence must be threat-specific, and all possible threats must be addressed as described in Chapter 22. Effective deterrence and mitigation then become a matter of strengthening the protections described throughout this chapter, which need not be very costly considering the response and recovery costs that could otherwise result. Premises or corporate security must deal with most threats of violence, but the infrastructure needs special protections to avoid disruption and to mitigate the downtime and cost consequences of any such event.

Reasonably good protection and mitigation measures can be simple and inexpensive. State and regional bomb squads or explosives units can advise and assist in many ways, including current briefings. Weapons of mass destruction (WMDs), other than nuclear weapons, are fast becoming a real threat, especially because many such devices are small and easily concealed. WMDs include chemical and biological agents and incendiary devices, while even small amounts of radioactive materials disbursed by an explosive *dirty bomb* can spread panic. The government considers these devices very serious threats, with businesses and their infrastructures as likely targets.

New federal office space must now be certified as bomb resistant, so that an explosion or the impact of a truck bomb cannot collapse the building. Officials can usually provide a current threat briefing and suggest protective measures. Although small areas may be destroyed, the structure will not collapse.

Just like some possible threats, there are also effective protection devices that cannot be mentioned. These are not marketed publicly and therefore are unknown to dealers, resellers, or distributors. The costs can be reasonable because they are only sold direct. Developers may supply classified systems for military or government use and offer declassified versions to other selected users. In this way, developers can restrict knowledge of their products to as few people as possible, so that others cannot discover how to recognize or circumvent them.

Many consultants who have worked with financial, regulated, or very large private companies know some of these specialty vendors. Usually, a consultant will approach

the vendor and discuss what is needed; the vendor may then contact the customer directly.

Finally, given today's environment of violence, get to know key local, state, and federal law enforcement and investigative officials, and ask their suggestions how best to protect an organization.

23.9.4 Medical Emergencies. Mitigating medical emergencies requires a first aid room on the premises, first aid and some medical supplies, a registered nurse if possible, and many workers trained in first aid and cardiopulmonary resuscitation (CPR). All security personnel and guards should be certified in first aid and CPR.

An automated external defibrillator (AED) on site will save lives and can be operated by anybody in an emergency. A portable AED currently costs about $1,000, and the suggested training is inexpensive. An AED is now required in all federally managed buildings. Many shopping malls, places of public assembly, and commercial aircraft are now equipped with one or more units.

Oxygen is often necessary to save lives and prevent permanent impairment. Most sites equipped with an AED also have oxygen units. Good portable units cost $800 or less.

23.10 COMPLETING THE SECURITY PLANNING PROCESS. The last step necessary to protect the information infrastructure has four components. Absent any of these components completed thoroughly, good security is not possible. They include:

1. Develop an all-hazard mitigation plan.
2. Develop all of the mitigation options for each identified threat and perform a cost-benefit analysis to determine which options are best.
3. Develop an overall security response plan to show who is responsible for what.
4. Complete the necessary implementation, accountability, and follow-up procedures.

23.10.1 All-Hazard Mitigation Plan. FEMA provides extensive resources for *Multi-Hazard Mitigation Planning* on a wide range of topics and including email updates.[13] *State and Local Mitigation Planning How-to Guides* can be helpful in this final step.[14]

- FEMA Publication 386-1, *Getting Started: Building Support for Mitigation Planning*, establishes the hazard-mitigation process.[15]
- Publication 386-2, *Understanding Your Risks: Identifying Hazards and Estimating Losses,* shows a method of cost-estimating potential losses due to flooding using tables rather than by calculation. These tables quickly show that losses can be far greater than expected.[16]
- Publication 386-3, *Developing the Mitigation Plan: Identifying Mitigation Actions and Implementation Strategies*, provides guidance on developing the mitigation strategy.[17]
- Publication 386-4, *Bringing the Plan to Life: Implementing the Hazard Mitigation Plan*, discusses guidance on implementation.[18]

- Publication 386-5, *Using Benefit-Cost Review in Mitigation Planning*, provides heuristics for selecting among mitigation options to maximize cost effectiveness of expenditures.[19]
- Publication 386-6, *Integrating Historic Property and Cultural Resource Considerations into Hazard Mitigation Planning*, discusses how to work with local planning authorities to protect cultural capital.[20]
- Publication 386-7, *Integrating Manmade Hazards Into Mitigation Planning*, explores how to plan for terrorist or criminal attacks in emergency planning.[21]
- Publication 386-8, *Multi-Jurisdictional Mitigation Planning*, emphasizes the importance of collaborative response planning (Federal Emergency Management Agency 2006).[22]

Once all possible threat situations have been identified and assessed as described in Chapter 22, the next step is to develop as many options as possible to prevent, deter, or mitigate disruption, injury, or damage from each threat. Although some threats cannot be prevented, there are always ways to prepare for and mitigate their impact. Usually, there are many mitigation options, so the objective is to determine which options are the most practical and affordable. The only objective way to do this is with cost-benefit analysis (described in Section 23.11.2).

In actuality, the options to protect against many different threats will be similar, but each option should be retained until the best mitigation strategy for each threat is determined. Here also there is likely to be one common mitigation strategy that covers many different threats. All credible threats should be listed in the mitigation plan, but the mitigation projects laid out will be far fewer in number.

When the next step is finished, the All-Hazard Mitigation Plan can be completed. The FEMA 386 how-to manuals described above provide the suggested format and content. FEMA's *How-To Guide #9* (FEMA 386-9), entitled *Using the Hazard Mitigation Plan to Prepare Successful Mitigation Projects*, is an extensive guide to using all the other resources mentioned above.[23]

The complete mitigation plan should be for official use only, and not released except to those with a need to know this information. The complete plan would be very helpful to a potential troublemaker because it shows where the organization is vulnerable, and to which threats. If disclosure of the complete plan is not well controlled, its contents could be leaked by, extorted from, or sold by an insider. However, the executive summary of the plan and abbreviated findings should be circulated widely, so that all stakeholders know that much is being done to protect them.

23.10.2 Cost-Benefit Analysis. Security is pointless unless it is cost effective and also adds value: that is, the cost of mitigating each threat must be less than the potential benefits and savings of the event not occurring, because if the protection is effective. The costs of every option can easily be determined; these are the initial and ongoing costs. Some future benefits, though, will be intangible, and all will have to be approximated. The long-term benefits of something not happening must reflect the approximated costs of:

- Disruptions that would reduce the productivity of the business and the performance of its information systems
- Morale and performance that could plummet, and remain very low because people feel unsafe and unprotected

- Loss of business or customers until operations could be restored to normal
- Response and recovery costs including extra time and overtime, expenses including lodging and meals, temporary facilities, public relations, and legal defense costs that are all likely to be incurred
- Legal, public relations, and other services and expenses to repair reputational damage, and fallen stock price, and to restore goodwill

Not all of these costs will follow every threat. But then again, there may also be additional, unexpected costs as well. In general, the response and recovery costs of any major security event tend to be far greater than expected. Nonetheless, each situation can be studied and some costs determined in order to facilitate a statistically valid cost-benefit analysis.

Cost and benefit information have no meaning unless each is associated with a common time frame. The likelihood of each threat should be assessed on an annualized basis (see Section 22.3.4), so that both the mitigation costs and the potential benefits can be amortized over the same life cycle.

There are many methods for cost-value analysis, and most are beyond the scope of this chapter. However, for those who are not financially trained, the federal government has a good system, freely available, that is widely required within the government. The U.S. government calls this system *BCA*, which stands for benefit-cost analysis. As part of this system, FEMA has developed a *Benefit-Cost Analysis Tool,* which assists grant applicants with financial analyses, such as net present value. The same site includes hazard data.[24]

One particular advantage of the BCA system is that the OMB publishes current cost-of-funds data needed to project any costs. Some of the private models tend to use wildly optimistic (or grossly out of date) future interest rates, which invalidates meaningful results. Again, it is wise to use a federal government model simply as a matter of risk avoidance. The BCA model is widely used and required for many grant applications.

23.10.3 Security Response Plan. A hazard mitigation plan is needed to document the threat assessment done in Chapter 22 and to list the mitigation options and the predicted costs and benefits associated with each; a security response plan is also needed to direct each stakeholder according to the type of threat experienced. The purpose of the security response plan is to define clearly who is in charge and who will do what, when and how, when any threat occurs.

The new NRF format revises and reestablishes the Emergency Support Functions (ESF) concept for each support activity that each organization may need.[25] There are now 16 numbered ESFs, beginning with ESF-1, "Transportation," ESF-2, "Communications & Alerting," and on through ESF-16, "Animal Health." It is recommended that the standard ESF titles be retained, even though many may not be applicable to a nongovernment organization, if only to maintain a uniformity and language that everyone understands. Additional ESFs will be needed to mount an effective response, but number these as 17 and upward.

One useful quick reference in the new response plan format is an Emergency Support Function Assignment Matrix (which is often Figure 1), a one-page graphic that shows which agency or department has primary, secondary, or support responsibility for each ESF. This is a handy quick-reference guide for management and staff when trouble comes.

Once a response plan is ready, standard operating procedures should be written by each affected department to outline the procedures it will use to respond. The term is in quotes because these are usually issued as guidelines so that there can be some flexibility to adjust to the actual conditions. The requirement for rigid adherence to a procedure is an invitation that invites litigation.

The complete security response plan should also be for official use only, and released only to those with a need to know. After all, this plan shows troublemakers just how the organization will respond. However, a press release summarizing the plan should be widely circulated so that all stakeholders know the organization is trying to protect them.

23.10.4 Implementation, Accountability, and Follow-Up. Once all the plans are completed and signed off by management and key officials, the job of implementation begins. Of the How-To Guides mentioned earlier, FEMA 386-4, *Bringing the Plan to Life,* will be helpful here.

This first implementation step is the most critical to the security-protection process. It begins with training so that everyone involved understands and accepts the procedures. There must then be periodic exercises and drills to test the response plan and to validate that the training has been effective. Plans that are not periodically tested are soon forgotten.

Every exercise and drill must be reviewed to determine what went right and, more important, what did not, and how to do better in the future. So too should every emergency response be reviewed. Documentation before, during, and after the event is important. There will be some lessons learned from each event, and these lessons should be used to update and improve the security systems to work better in the future.

It is also critical to establish accountability for the infrastructure security. There should be only one person in charge of each function. Responsibility cannot be spread among management or departments, nor can responsibility be worn as a second hat for someone with many other duties. The senior responsible authority must set up schedules for periodic review and update of the plans and procedures, training, and exercises to make sure that the security program remains current and effective.

Good security management must also include oversight. Good planning must begin in the boardroom, and the directors must also provide continuing oversight to ascertain good security. Outside, independent, auditors who are directed from the boardroom are best able to validate the current condition of the security program. The auditor's written opinion is evidence of whether the organization is fully compliant or not. The best procedure for an infrastructure security audit is suggested in Section 23.2.7.

For management's own peace of mind (and possibly as a requirement for maintaining insurance and obtaining credit), there may be periodic security inspections, testing of defenses, and some penetration tests, including deceptions to gain access done by independent professionals.

23.11 FURTHER READING

The further readings suggested in Chapter 22 will also be helpful in this chapter.

The Congressional Research Service report to Congress dated June 2006, entitled *Federal Emergency Management and Homeland Security Organization: Historical Development and Legislative Options*, is interesting reading for anyone burdened with security planning and management. This outlines the federal government's struggles since 1947 to get the system right. It also lists legislation

before the 109th Congress to fix the problems. The report can be downloaded free at www.fas.org/sgp/crs/homesec/RL33369.pdf

For an interesting discussion of the inner workings of federal response systems and how better security planning and management can be implemented, see Christopher Cooper and Robert Block, *Disaster: Hurricane Katrina and the Failure of Homeland Security*, Holt Paperbacks (ISBN 978-0805086508), 2007: 352 pp.

23.12 NOTES

1. Federal Emergency Management Agency (FEMA), "National Incident Management System," FEMA Website, January 30, 2013, www.fema.gov/national-incident-management-system (Federal Emergency Management Agency 2013).
2. FEMA, "National Response Framework," FEMA Website, August 20, 2012, www.fema.gov/national-response-framework (Federal Emergency Management Agency 2012).
3. FEMA, "IS-860 National Infrastructure Protection Plan (NIPP) Course," FEMA Emergency Management Institute, May 24, 2011, www.training.fema.gov/EMIWeb/IS/is860.asp (Federal Emergency Management Agency 2011).
4. FEMA, "Hazard Mitigation Planning Resources," FEMA Website, June 15, 2012, www.fema.gov/hazard-mitigation-planning-resources
5. "Welcome to Sarbanes Oxley 101," Sarbanes-Oxley-101.com Website, January, 2013, www.sarbanes-oxley-101.com
6. Department of Health & Human Services, "Health Information Privacy," hhs.gov Website, January 2013, www.hhs.gov/ocr/privacy/hipaa/understanding/index.html
7. Federal Trade Commission, "How To Comply with the Privacy of Consumer Financial Information Rule of the Gramm-Leach-Bliley Act." Bureau of Consumer Protection Business Center, July 2002, http://business.ftc.gov/documents/bus67-how-comply-privacy-consumer-financial-information-rule-gramm-leach-bliley-act
8. Emergency Management Accreditation Program, "EMAP Home." EMAP Website, 2010, www.emaponline.org/index.php
9. Maguire, Valerie, "TIA Publishes New Standard for Telecommunication Administration," Network Infrastructure Blog, July 20, 2012, http://blog.siemon.com/infrastructure/tia-publishes-new-standard-for-telecommunication-administration
10. National Fire Protection Association, "NFPA 70: National Electrical Code—Current Edition 2011," NFPA Codes & Standards, 2011, www.nfpa.org/aboutthecodes/AboutTheCodes.asp?DocNum=70
11. Ward, Mike, "Start of airport-style security at Texas Capitol: Some visitors complain that licensed gun holders are sped through," Austin American Statesman, May 21, 2010, www.statesman.com/news/news/state-regional-govt-politics/start-of-airport-style-security-at-texas-capitol/nRs3 C
12. Department of Homeland Security, "Airport Passenger Screening: Backscatter X-Ray Machines: Solicitation Number RSEN-13-00004," Federal Business Opportunities Website, December 13, 2012, www.fbo.gov/index?s=opportunity&mode=form&id=0af3059baf5a4b278a75936152e93253&tab=core&_cview=0
13. FEMA, "Multi-Hazard Mitigation Planning," FEMA Website, September 12, 2012, www.fema.gov/multi-hazard-mitigation-planning
14. FEMA, "State and Local Mitigation Planning Fact Sheet," FEMA Website, May 2011, www.fema.gov/pdf/media/factsheets/2011/mit_state_local_plan.pdf

15. FEMA, "Mitigation Planning How-To Guide #1: Getting Started: Building Support for Mitigation Planning (FEMA 386-1)," FEMA Website, September 2002, www.fema.gov/library/viewRecord.do?id=1867
16. FEMA, "Mitigation Planning How-To Guide #2: Understanding Your Risks: Identifying Hazards and Estimating Losses (FEMA 386-2)," FEMA Website, August 2001, www.fema.gov/library/viewRecord.do?id=1880
17. FEMA, "Mitigation Planning How-To Guide #3: Developing the Mitigation Plan: Identifying Mitigation Actions and Implementation Strategies (FEMA 386-3)," FEMA Website, April 2003, www.fema.gov/library/viewRecord.do?id=1886
18. FEMA, "Mitigation Planning How-To Guide #4: Bringing the Plan to Life: Implementing the Hazard Mitigation Plan (FEMA 386-4)," FEMA Website, August 2003, www.fema.gov/library/viewRecord.do?id=1887
19. FEMA, "Mitigation Planning How-To Guide #5: Using Benefit-Cost Review in Mitigation Planning (FEMA 386-5)," FEMA Website, May 2007, www.fema.gov/library/viewRecord.do?id=2680
20. FEMA, "Integrating Historic Property and Cultural Resource Considerations into Hazard Mitigation Planning (FEMA 386-6)," FEMA Website, May 2005, www.fema.gov/library/viewRecord.do?id=1892
21. FEMA, "Integrating Manmade Hazards Into Mitigation Planning (FEMA 386-7)." FEMA Website, September 2003, www.fema.gov/library/viewRecord.do?id=1915
22. FEMA, "Multi-Jurisdictional Mitigation Planning (FEMA 386-8)," FEMA Website, August 2006, www.fema.gov/library/viewRecord.do?id=1905
23. FEMA, "Mitigation Planning How-To #9: Using the Hazard Mitigation Plan to Prepare Successful Mitigation Projects (FEMA 386-9)," FEMA Website, August 2008, www.fema.gov/library/viewRecord.do?id=3388
24. FEMA, "Benefit-Cost Analysis," FEMA Website, February 1, 2013, www.fema.gov/benefit-cost-analysis
25. FEMA, "National Response Network," FEMA Website, August 20, 2012, www.fema.gov/national-response-framework

CHAPTER 24

OPERATING SYSTEM SECURITY

William Stallings

24.1 INFORMATION PROTECTION AND SECURITY **24·1**

24.2 REQUIREMENTS FOR OPERATING SYSTEM SECURITY **24·2**

24.2.1 Requirements 24·2
24.2.2 Computer System Assets 24·3
24.2.3 Design Principles 24·4

24.3 PROTECTION MECHANISMS **24·4**

24.3.1 Protection of Memory 24·5
24.3.2 User-Oriented Access Control 24·6
24.3.3 Data-Oriented Access Control 24·7
24.3.4 Protection Based on an Operating System Mode 24·9
24.3.5 Protection Based on Virtualization 24·10

24.4 FILE SHARING **24·11**

24.4.1 Access Rights 24·11
24.4.2 Simultaneous Access 24·12

24.5 TRUSTED SYSTEMS **24·12**

24.5.1 Trojan Horse Defense 24·14

24.6 WINDOWS 2000 SECURITY **24·14**

24.6.1 Access-Control Scheme 24·16
24.6.2 Access Token 24·16
24.6.3 Security Descriptors 24·17

24.7 FURTHER READING **24·20**

24.8 NOTES **24·21**

24.1 INFORMATION PROTECTION AND SECURITY. This chapter reviews the principles of security in operating systems. Some general-purpose tools can be built into computers and operating systems (OSs) that support a variety of protection and security mechanisms. In general, the concern is with the problem of controlling access to computer systems and the information stored in them.

One of the core concepts in all such discussions is the *process*, which is defined as the execution of a specific piece of code by a particular user at a particular time on a particular processor.

Four types of overall protection policies, of increasing order of difficulty, have been identified:

1. **No sharing.** In this case, processes are completely isolated from each other, and each process has exclusive control over the resources statically or dynamically assigned to it. With this policy, processes often "share" a program or data file by making a copy of it and transferring the copy into their own virtual memory.

2. **Sharing originals of program or data files.** With the use of reentrant code, a single physical realization of a program can appear in multiple virtual address spaces, as can read-only data files. Special locking mechanisms are required for the sharing of writable data files, to prevent simultaneous users from interfering with each other.
3. **Confined, or memoryless, subsystems.** In this case, processes are grouped into subsystems to enforce a particular protection policy. For example, a "client" process calls a "server" process to perform some task on data. The server is to be protected against the client discovering the algorithm by which it performs the task, while the client is to be protected against the server's retaining any information about the task being performed.
4. **Controlled information dissemination.** In some systems, security classes are defined to enforce a particular dissemination policy. Users and applications are given security clearances of a certain level, while data and other resources (e.g., input/output [I/O] devices) are given security classifications. The security policy enforces restrictions concerning which users have access to which classifications. This model is useful not only in the military context but in commercial applications as well.[1]

Much of the work in security and protection as it relates to OSs can be roughly grouped into three categories.

1. **Access control.** Concerned with regulating user access to the total system, subsystems, and data, and regulating process access to various resources and objects within the system.
2. **Information flow control.** Regulates the flow of data within the system and its delivery to users.
3. **Certification.** Relates to proving that access and flow control mechanisms perform according to their specifications and that they enforce desired protection and security policies.

This chapter looks at some of the key mechanisms for providing OS security and then examines Windows 2000 as a case study.

24.2 REQUIREMENTS FOR OPERATING SYSTEM SECURITY

24.2.1 Requirements. Understanding the types of threats to OS security requires a definition of security requirements. OS security addresses four requirements:

1. **Confidentiality**. Requires that the information in a computer system be accessible only for reading by authorized parties. This type of access includes printing, displaying, and other forms of disclosure, including simply revealing the existence of an object.
2. **Integrity**. Requires that only authorized parties be able to modify computer system assets. Modification includes writing, changing, changing status, deleting, and creating.

3. **Availability.** Requires that computer system assets are available to authorized parties.
4. **Authentication.** Requires that a computer system be able to verify the identifier of a user, a device or a process.

24.2.2 Computer System Assets. The assets of a computer system can be categorized as hardware, software, and data.

24.2.2.1 Hardware. The main threat to computer system hardware is in the area of availability. Hardware is the most vulnerable to attack and the least amenable to automated controls. Threats include accidental and deliberate damage to equipment as well as theft. The proliferation of personal computers and workstations and the increasing use of local area networks (LANs) increase the potential for losses in this area. Physical and administrative security measures are needed to deal with these threats. Chapters 22 and 23 in this *Handbook* discuss physical security.

24.2.2.2 Software. The OS, utilities, and application programs are what make computer system hardware useful to businesses and individuals. Several distinct threats need to be considered.

A key threat to software is an attack on availability. Software, especially application software, is surprisingly easy to delete. Software also can be altered or damaged to render it useless or dangerous. Careful software configuration management, which includes making backups of the most recent version of software, can maintain high availability. A more difficult problem to deal with is software modification that results in a program that still functions but that behaves differently from before. A final problem is control or possession of software. Although certain countermeasures are available, by and large the problem of unauthorized copying of software has not been solved.

Chapters 38, 39, and 40 of this *Handbook* discuss software security in some detail.

24.2.2.3 Data. Hardware and software security typically are concerns of computing center professionals or individual concerns of personal computer users. A much more widespread problem is data security, which involves files and other forms of data controlled by individuals, groups, and business organizations.

Security concerns with respect to data are broad, encompassing confidentiality, control or possession, integrity, authenticity, availability and utility. For a sound theoretical treatment of the attributes of information that must be protected through security measures, see Chapter 3 in this *Handbook*.

In the case of availability, the concern is with the *destruction* of data files, which can occur either accidentally or maliciously, and with delays in timely access to data.

The obvious concern with confidentiality is the unauthorized reading of data files or databases, and this area has been the subject of perhaps more research and effort than any other area of computer security. A less obvious secrecy threat involves the analysis of data and manifests itself in the use of so-called statistical databases or *data mining*, which provide summary or aggregate information and potentially lead to discovery of unpublicized tendencies, relations, or trends. Aggregate information does not necessarily threaten the privacy of the individuals involved. However, as the use of statistical databases grows, there is an increasing potential for disclosure of personal information through induction or deduction. In essence, characteristics of constituent

individuals may be identified through careful analysis. To take a simple example, if one table records the aggregate of the incomes of respondents A, B, C, and D and another records the aggregate of the incomes of A, B, C, D, and E, the difference between the two aggregates would be the income of E. This problem is exacerbated by the increasing desire to combine data sets. In many cases, matching several sets of data for consistency at levels of aggregation appropriate to the problem requires a retreat to elemental units in the process of constructing the necessary aggregates. Thus, the elemental units, which are the subject of privacy concerns, are available at various stages in the processing of data sets.

Finally, data integrity is a major concern in most installations. Modifications to data files can have consequences ranging from minor to disastrous.

24.2.3 Design Principles. Saltzer and Schroeder identify a number of principles for the design of security measures for the various threats to computer systems. These include:

- **Least privilege**. Every program and every user of the system should operate using the least set of privileges necessary to complete the job. Access rights should be acquired by explicit permission only; the default should be "no access."
- **Economy of mechanisms**. Security mechanisms should be as small and simple as possible, aiding in their verification. This usually means that they must be an integral part of the design rather than add-on mechanisms to existing designs.
- **Acceptability**. Security mechanisms should not interfere unduly with the work of users. At the same time, the mechanisms should meet the needs of those who authorize access. If the mechanisms are not easy to use, they are likely to be unused or incorrectly used.
- **Complete mediation**. Every access must be checked against the access-control information, including those accesses occurring outside normal operation, as in recovery or maintenance.
- **Open design**. The security of the system should not depend on keeping the design of its mechanisms secret. Thus, the mechanisms can be reviewed by many experts, and users can have high confidence in them.[2]

24.3 PROTECTION MECHANISMS.

The introduction of multiprogramming brought about the ability to share resources among users. This sharing involves not just the processor but also:

- Memory
- I/O devices, such as disks and printers
- Programs
- Data

The ability to share these resources introduced the need for protection. Pfleeger and Pfleeger point out that an OS may offer protection along this spectrum:

- **No protection**. This is appropriate when sensitive procedures are being run at separate times.

- **Isolation**. This approach implies that each process operates separately from other processes, with no sharing or communication. Each process has its own address space, files, and other objects.
- **Share all or share nothing**. The owner of an object (e.g., a file or memory segment) declares it to be public or private. In the former case, any process may access the object; in the latter, only the owner's processes may access the object.
- **Share via access limitation**. The OS checks the permissibility of each access by a specific user to a specific object. The OS therefore acts as a guard, or gatekeeper, between users and objects, ensuring that only authorized accesses occur.
- **Share via dynamic capabilities**. This extends the concept of access control to allow dynamic creation of sharing rights for objects.
- **Limit use of an object**. This form of protection limits not just access to an object but the use to which that object may be put. For example, a user may be allowed to view a sensitive document but not print it. Another example is that a user may be allowed access to a database to derive statistical summaries but not to determine specific data values.[3]

The preceding items are listed roughly in increasing order of difficulty to implement but also in increasing order of fineness of protection that they provide. A given OS may provide different degrees of protection for different objects, users, or applications.

The OS needs to balance the need to allow sharing, which enhances the utility of the computer system, with the need to protect the resources of individual users. This section considers some of the mechanisms by which OSs have enforced protection for these objects.

24.3.1 Protection of Memory. In a multiprogramming environment, protection of main memory (random-access memory, or *RAM*) is essential. The concern here is not just security but the correct functioning of the various processes that are active. If one process can inadvertently write into the memory space of another process, then the latter process may not execute properly.

The separation of the memory space of various processes is accomplished easily with a virtual memory scheme. Either segmentation or paging, or the two in combination, provides an effective means of managing main memory. If complete isolation is sought, then the OS simply must ensure that each segment or page is accessible only by the process to which it is assigned. This is accomplished easily by requiring that there be no duplicate entries in page and/or segment tables.

If sharing is to be allowed, then the same segment or page may appear in more than one table. This type of sharing is accomplished most easily in a system that supports segmentation or a combination of segmentation and paging. In this case, the segment structure is visible to the application, and the application can declare individual segments to be sharable or nonsharable. In a pure paging environment, it becomes more difficult to discriminate between the two types of memory, because the memory structure is transparent to the application.

Segmentation, especially, lends itself to the implementation of protection and sharing policies. Because each segment table entry includes a length as well as a base address, a program cannot inadvertently access a main memory location beyond the limits of a segment. To achieve sharing, it is possible for a segment to be referenced in the segment tables of more than one process. The same mechanisms are available in a paging system. However, in this case the page structure of programs and data is not visible

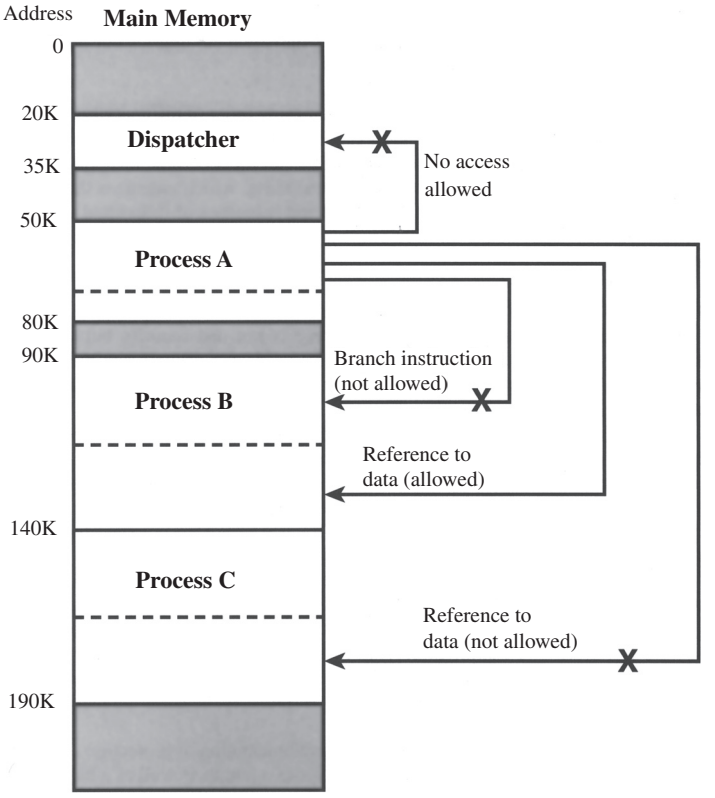


EXHIBIT 24.1 Protection Relationships between Segments

to the programmer, making the specification of protection and sharing requirements more awkward. Exhibit 24.1 illustrates the types of protection relationships that can be enforced in such a system.

An example of the hardware support that can be provided for memory protection is that of the IBM System/370 family of machines, on which OS/390 runs. Associated with each page frame in main memory is a 7-bit storage control key, which may be set by the OS. Two of the bits indicate whether the page occupying this frame has been referenced and changed; these bits are used by the page replacement algorithm. The remaining bits are used by the protection mechanism: a 4-bit access-control key and a fetch-protection bit. Processor references to memory and direct memory access (DMA). DMA I/O memory references must use a matching key to gain permission to access that page. The fetch-protection bit indicates whether the access-control key applies to writes or to both reads and writes. In the processor, there is a program status word (PSW), which contains control information relating to the process that is currently executing. Included in this word is a 4-bit PSW key. When a process attempts to access a page or to initiate a DMA operation on a page, the current PSW key is compared to the access code. A write operation is permitted only if the codes match. If the fetch bit is set, then the PSW key must match the access code for read operations.

24.3.2 User-Oriented Access Control. The measures taken to control access in a data processing system fall into two categories: those associated with the user and those associated with the data.

The most common technique for user access control on a shared system or server is the user logon, which requires both a user identifier (ID) and some form of *authentication*, such as providing a password, a token, or biometric attributes. Authentication refers to the binding of a real-world identity (for example, a named employee or a named role in an organization) and the ID being presented. The system will allow a user to log on only if that user's ID is known to the system and if the user knows the password associated by the system with that ID.

Once a user has established a *session*, the operating system can then *authorize* different forms of access (e.g., read, write, append, lock, or execute) to different types of data (e.g., specific files, databases, devices, or communications).

User access control in a distributed environment can be either centralized or decentralized. In a centralized approach, the network provides a logon service, determining who is allowed to use the network and to whom the user is allowed to connect.

Decentralized user access control treats the network as a transparent communication link, and the destination host carries out the usual logon procedure. The security concerns for transmitting passwords over the network must still be addressed.

In many networks, two levels of access control may be used. Individual hosts may be provided with a logon facility to protect host-specific resources and application. In addition, the network as a whole may provide protection to restrict network access to authorized users. This two-level facility is desirable for the common case, currently, in which the network connects disparate hosts and simply provides a convenient means of terminal-host access. In a more uniform network of hosts, some centralized access policy could be enforced in a network control center.

Chapters 28 and 29 of this *Handbook* present more information about identification and authentication.

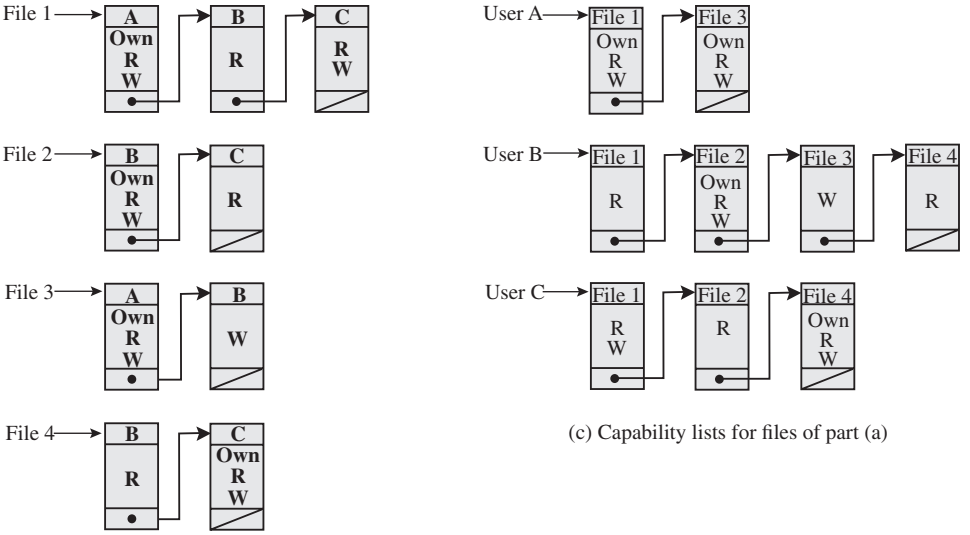
24.3.3 Data-Oriented Access Control. Following successful logon, the user is granted access to one or a set of hosts and applications. This is generally not sufficient for a system that includes sensitive data in its database. Through the user access-control procedure, a user can be identified to the system. Associated with each user, there can be a profile that specifies permissible operations and file accesses. The OS can then enforce rules based on the user profile. The database-management system, however, must control access to specific records or even portions of records. For example, it may be permissible for anyone in administration to obtain a list of company personnel, but only selected individuals may have access to salary information. The issue is more than just one of level of detail. Whereas the OS may grant a user permission to access a file or use an application, following which there are no further security checks, the database management system must make a decision on each individual access attempt. That decision will depend not only on the user's identity but also on the specific parts of the data being accessed and even on the information already divulged to the user.

A general model of access control as exercised by a file or database management system is that of an *access matrix* (see Exhibit 24.2a). The basic elements of the model are:

- **Subject**. An entity capable of accessing objects. Generally, the concept of subject equates with that of process. Any user or application actually gains access to an object by means of a process that represents that user or application.
- **Object**. Anything to which access is controlled. Examples include files, portions of files, programs, and segments of memory.
- **Access right**. The way in which an object is accessed by a subject. Examples are read, write, and execute.

	File 1	File 2	File 3	File 4	Account 1	Account 2
User A	Own R W		Own R W		Inquiry Credit	
User B	R	Own R W	W	R	Inquiry Debit	Inquiry Credit
User C	R W	R		Own R W		Inquiry Debit

(a) Access matrix



(b) Access control lists for files of part (a)

(c) Capability lists for files of part (a)

EXHIBIT 24.2 Example of Access-Control Structures
Source: Based on a figure in Sandhu (1996).

One dimension of the matrix consists of identified subjects that may attempt data access. Typically, this list will consist of individual users or user groups, although access could be controlled for terminals, hosts, or applications, instead of, or in addition to, users. The other dimension lists the objects that may be accessed. At the greatest level of detail, objects may be individual data fields. More aggregate groupings, such as records, files, or even the entire database, also may be objects in the matrix. Each entry in the matrix indicates the access rights of that subject for that object.

In practice, an access matrix usually is sparse and is implemented by decomposition in one of two ways. The matrix may be decomposed by columns, yielding *access-control lists* (see Exhibit 24.2b). Thus, for each object, an access-control list lists users and their permitted access rights. The access-control list may contain a default, or public, entry. This allows users who are not explicitly listed as having special rights to have a default set of rights. Elements of the list may include individual users as well as groups of users.

Decomposition by rows yields *capability tickets* (see Exhibit 24.2c). A capability ticket specifies authorized objects and operations for a user. Each user has a number

of tickets and may be authorized to lend or give them to others. Because tickets may be dispersed around the system, they present a greater security problem than access-control lists. In particular, the ticket must be unforgeable. One way to accomplish this is to have the OS hold all tickets on behalf of users. These tickets would have to be held in a region of memory inaccessible to users.

Network considerations for data-oriented access control parallel those for user-oriented access control. If only certain users are permitted to access certain items of data, then encryption may be needed to protect those items during transmission to authorized users. Typically, data access control is decentralized, that is, controlled by host-based database management systems. If a network database server exists on a network, then data access control becomes a network function.

24.3.4 Protection Based on an Operating System Mode. One technique used in all OSs to provide protection is based on the mode of processor execution. Most processors support at least two modes of execution: the mode normally associated with the OS and that normally associated with user programs. Certain instructions can be executed only in the more privileged mode. These would include reading or altering a control register, such as the program status word; primitive I/O instructions; and instructions that relate to memory management. In addition, certain regions of memory can be accessed only in the more privileged mode.

The less privileged mode often is referred to as the *user* mode, because user programs typically would execute in this mode. The more privileged mode is referred to as the *system mode*, *control mode*, or *kernel mode*. This last term refers to the kernel of the OS, which is that portion of the OS that encompasses the important system functions. Exhibit 24.3 lists the functions typically found in the kernel of an OS.

EXHIBIT 24.3 Typical Kernel Mode Operating System Functions

Process Management

- Process creation and termination
- Process scheduling and dispatching
- Process switching
- Process synchronization and support for interprocess communication
- Management of process control blocks

Memory Management

- Allocation of address space to processes
- Swapping
- Page and segment management

I/O Management

- Buffer management
- Allocation of I/O channels and devices to processes

Support functions

- Interrupt handling
- Accounting
- Monitoring

The reason for using two modes should be clear. It is necessary to protect the OS and key OS tables, such as process control blocks, from interference by user programs. In the kernel mode, the software has complete control of the processor and all its instructions, registers, and memory. This level of control is not necessary, and for safety is not desirable, for user programs.

Two questions arise: How does the processor know in which mode it is to be executing, and how is the mode changed? Regarding the first question, typically there is a bit in the program status word that indicates the mode of execution. This bit is changed in response to certain events. For example, when a user makes a call to an OS service, the mode is set to the kernel mode. Typically, this is done by executing an instruction that changes the mode. When the user makes a system service call, or when an interrupt transfers control to a system routine, the routine executes the change-mode instruction to enter a more privileged mode and executes it again to enter a less privileged mode before returning control to the user process. If a user program attempts to execute a change-mode instruction, it will simply result in a call to the OS, which will return an error unless the mode change is to be allowed.

More sophisticated mechanisms also can be provided. A common scheme is to use a ring-protection structure. In this scheme, lower-numbered, or inner, rings enjoy greater privilege than higher-numbered, or outer, rings. Typically, ring 0 is reserved for kernel functions of the OS, with applications at a higher level. Some utilities or OS services may occupy an intermediate ring. Basic principles of the ring system are:

- A program may access only those data that reside on the same ring or a less privileged ring.
- A program may call services residing on the same or a more privileged ring.

An example of the ring protection approach is found on the VAX VMS OS, which uses four modes:

1. **Kernel**. Executes the kernel of the VMS OS, which includes memory management, interrupt handling, and I/O operations.
2. **Executive**. Executes many of the OS service calls, including file and record (disk and tape) management routines.
3. **Supervisor**. Executes other OS services, such as responses to user commands.
4. **User**. Executes user programs, plus utilities such as compilers, editors, linkers, and debuggers.

A process executing in a less privileged mode often needs to call a procedure that executes in a more privileged mode; for example, a user program requires an OS service. This call is achieved by using a change-mode (CHM) instruction, which causes an interrupt that transfers control to a routine at the new access mode. A return is made by executing the REI (return from exception or interrupt) instruction.

24.3.5 Protection Based on Virtualization. With the growing availability of memory (e.g., tens to hundreds of gigabytes of RAM), disk space (terabytes to petabytes of storage), faster processors (tens of gigahertz), and multicore systems (potentially thousands of processors working in parallel), *virtualization* of computers has progressed to practical and widespread usability. Instantiations of an operating environment can coexist using shared resources without permitting any direct communication among them. Each instantiation is encapsulated and completely protected against

intrusion or interference from processes running on other virtual machines sharing the same physical resources.

24.4 FILE SHARING. Multiuser systems almost always require that files can be shared among a number of users. Two issues arise: access rights and the management of simultaneous access.

24.4.1 Access Rights. The file system should provide a flexible tool for allowing extensive file sharing among users. The file system should provide a number of options so that the way in which a particular file is accessed can be controlled. Typically, users or groups of users are granted certain access rights to a file. A wide range of access rights has been used. The next list indicates access rights that can be assigned to a particular user for a particular file.

- **None**. The user may not even learn of the existence of the file, much less access it. To enforce this restriction, the user would not be allowed to read the user directory that includes this file.
- **Knowledge**. The user can determine that the file exists and who its owner is. The user is then able to petition the owner for additional access rights.
- **Execution**. The user can load and execute a program but cannot copy it. Proprietary programs often are made accessible with this restriction.
- **Locking**. The user can change the status of a logical *flag* that indicates temporary restrictions on access to data. Database management systems provide for locking to control concurrent access to records so that different processes can avoid overwriting each other's modifications.
- **Reading**. The user can read the file for any purpose, including copying and execution. Some systems are able to enforce a distinction between viewing and copying. In the former case, the contents of the file can be displayed to the user, but the user has no means for making a copy.
- **Appending**. The user can add data to the file, often only at the end, but cannot modify or delete any of the file's contents. This right is useful in collecting data from a number of sources.
- **Updating**. The user can modify, delete, and add to the file's data. This normally includes writing the file initially, rewriting it completely or in part, and removing all or a portion of the data. Some systems distinguish among different degrees of updating.
- **Changing protection**. The user can change the access rights granted to other users. Typically, only the owner of the file holds this right. In some systems, the owner can extend this right to others. To prevent abuse of this mechanism, the file owner typically is able to specify which rights can be changed by the holder of this extended right.
- **Deletion**. The user can delete the file from the file system.

These rights can be considered to constitute a hierarchy, with each right implying those that precede it. Thus, if a particular user is granted the updating right for a particular file, then that user also is granted these rights: knowledge, execution, reading, and appending.

One user is designated as owner of a given file, usually this is the person who initially created the file. The owner has all of the access rights listed previously and may grant rights to others. Access can be provided to different classes of users:

- **Specific user**. Individual users who are designated by user ID.
- **User groups**. A set of users who are not individually defined. The system must have some way of keeping track of the membership of user groups.
- **All**. All users who have access to this system. These are public files.

24.4.2 Simultaneous Access. When access is granted to append or update a file to more than one user, the OS or file management system must enforce discipline. A brute-force approach is to allow a user to lock the entire file when it is to be updated. A finer grain of control is to lock individual records during update. Issues of mutual exclusion and deadlock must be addressed in designing the shared access capability.

Chapter 52 in this *Handbook* discusses application controls in more detail.

24.5 TRUSTED SYSTEMS. Much of what has been discussed so far has concerned protecting a given message or item from passive or active attack by a given user. A somewhat different but widely applicable requirement is to protect data or resources on the basis of levels of security. This is commonly found in the military, where information is categorized as unclassified (U), confidential (C), secret (S), top secret (TS), or beyond. This concept is equally applicable in other areas, where information can be organized into gross categories and users can be granted clearances to access certain categories of data. For example, the highest level of security might be for strategic corporate planning documents and data, accessible only by corporate officers and their staff; next might come sensitive financial and personnel data, accessible only by administration personnel, corporate officers, and so on.

When multiple categories or levels of data are defined, the requirement is referred to as *multilevel security.* The general statement of the requirement for multilevel security is that a subject at a high level may not convey information to a subject at a lower or incomparable level unless that flow accurately reflects the will of an authorized user. For implementation purposes, this requirement is in two parts and is simply stated. A multilevel secure system must enforce:

1. **No read up**. A subject can only read an object of less or equal security level. This is referred to in the literature as the *simple security property.*
2. **No write down**. A subject can only write into an object of greater or equal security level. This is referred to in the literature as the **-property* (pronounced *star property*).

These two rules, if properly enforced, provide multilevel security. For a data processing system, the approach that has been taken, and has been the object of much research and development, is based on the *reference monitor* concept. This approach is depicted in Exhibit 24.4. The reference monitor is a controlling element in the hardware and OS of a computer that regulates the access of subjects to objects on the basis of security parameters of the subject and object. The reference monitor has access to a file, known as the *security kernel database*, that lists the access privileges (security clearance) of each subject and the protection attributes (classification level) of each

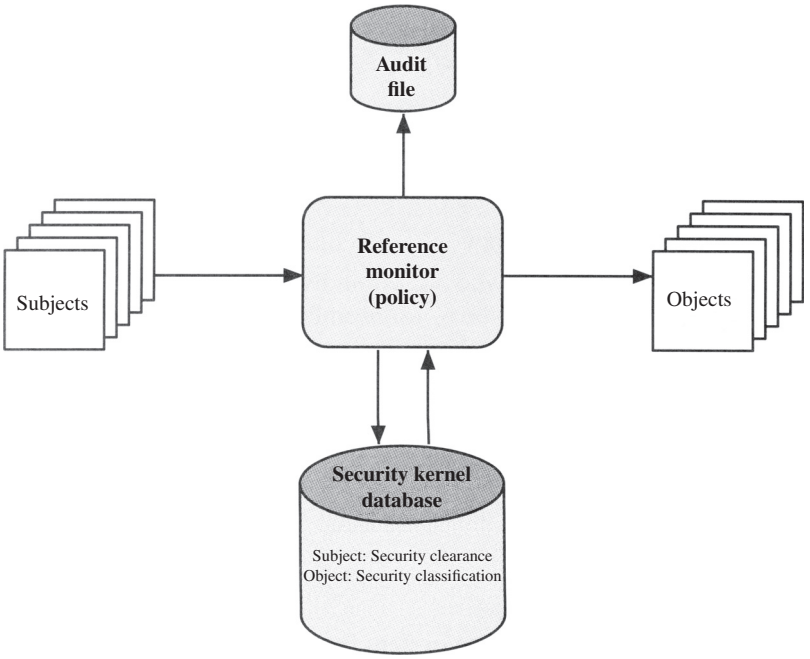


EXHIBIT 24.4 Reference Monitor Concept

object. The reference monitor enforces the security rules (no read up, no write down) and has these properties:

- **Complete mediation**. The security rules are enforced on every access, not just, for example, when a file is opened.
- **Isolation**. The reference monitor and database are protected from unauthorized modification.
- **Verifiability**. The reference monitor's correctness must be provable. That is, it must be possible to demonstrate mathematically that the reference monitor enforces the security rules and provides complete mediation and isolation.

These are stiff requirements. The requirement for complete mediation means that every access to data within main memory and on disk and tape must be mediated. Pure software implementations impose too high a performance penalty to be practical; the solution must be at least partly in hardware. The requirement for isolation means that it must not be possible for an attacker, no matter how clever, to change the logic of the reference monitor or the contents of the security kernel database. Finally, the requirement for mathematical proof is formidable for something as complex as a general-purpose computer. A system that can provide such verification is referred to as a *trusted system.*

Chapter 9 in this *Handbook* discusses mathematical models of computer security in more detail.

A final element illustrated in Exhibit 24.4 is an audit file. Important security events, such as detected security violations and authorized changes to the security kernel database, are stored in the audit file.

In an effort to meet its own needs and as a service to the public, the U.S. Department of Defense in 1981 established the Computer Security Center within the National Security Agency (NSA), with the goal of encouraging the widespread availability of trusted computer systems. This goal is realized through the center's Commercial Product Evaluation Program. In essence, the center attempts to evaluate commercially available products as meeting the security requirements just outlined. The center classifies evaluated products according to the range of security features that they provide. These evaluations are needed for Department of Defense procurements but are published and freely available. Hence, they can serve as guidance to commercial customers for the purchase of commercially available, off-the-shelf equipment.

24.5.1 Trojan Horse Defense. A *Trojan horse* attack involves software that appears to have acceptable functions but which conceals additional, unauthorized functionality. Chapters 2 and 15 in this *Handbook* provide details of many Trojan horse attacks.

One way to secure against Trojan horse attacks is by the use of a secure, trusted OS. Exhibit 24.5 illustrates an example. In this case, a Trojan horse is used to get around the standard security mechanism used by most file management and OSs: the access-control list. In this example, a user named Bob interacts through a program with a data file containing the critically sensitive character string "CPE170KS". User Bob has created the file with read/write permission provided only to programs executing on his own behalf: that is, only processes that are owned by Bob may access the file.

The Trojan horse attack begins when a hostile user, named Alice, gains legitimate access to the system and installs both a Trojan horse program and a private file to be used in the attack as a "back pocket." Alice gives read/write permission to herself for this file and gives Bob write-only permission (see Exhibit 24.5a). Alice now induces Bob to invoke the Trojan horse program, perhaps by advertising it as a useful utility. When the program detects that it is being executed by Bob, it reads the sensitive character string from Bob's file and copies it into Alice's back-pocket file (see Exhibit 24.5b). Both the read and write operations satisfy the constraints imposed by access-control lists. Alice then has only to access Bob's file at a later time to learn the value of the string.

Now consider the use of a secure OS in this scenario (see Exhibit 24.5c). Security levels are assigned to subjects at logon on the basis of criteria such as the terminal from which the computer is being accessed and the user involved, as identified by password/ID. In this example, there are two security levels, sensitive (gray) and public (white), ordered so that sensitive is higher than public. Processes owned by Bob and Bob's data file are assigned the security level sensitive. Alice's file and processes are restricted to public. If Bob invokes the Trojan horse program (see Exhibit 24.5d), that program acquires Bob's security level. It is therefore able, under the simple security property, to observe the sensitive character string. When the program attempts to store the string in a public file (the back-pocket file), however, the *-property is violated and the attempt is disallowed by the reference monitor. Thus, the attempt to write into the back-pocket file is denied even though the access-control list permits it: The security policy takes precedence over the access-control list mechanism.

24.6 WINDOWS 2000 SECURITY. A good example of the access-control concepts discussed in this chapter is the Windows 2000 (W2K) access-control facility, which exploits object-oriented concepts to provide a powerful and flexible access-control capability.

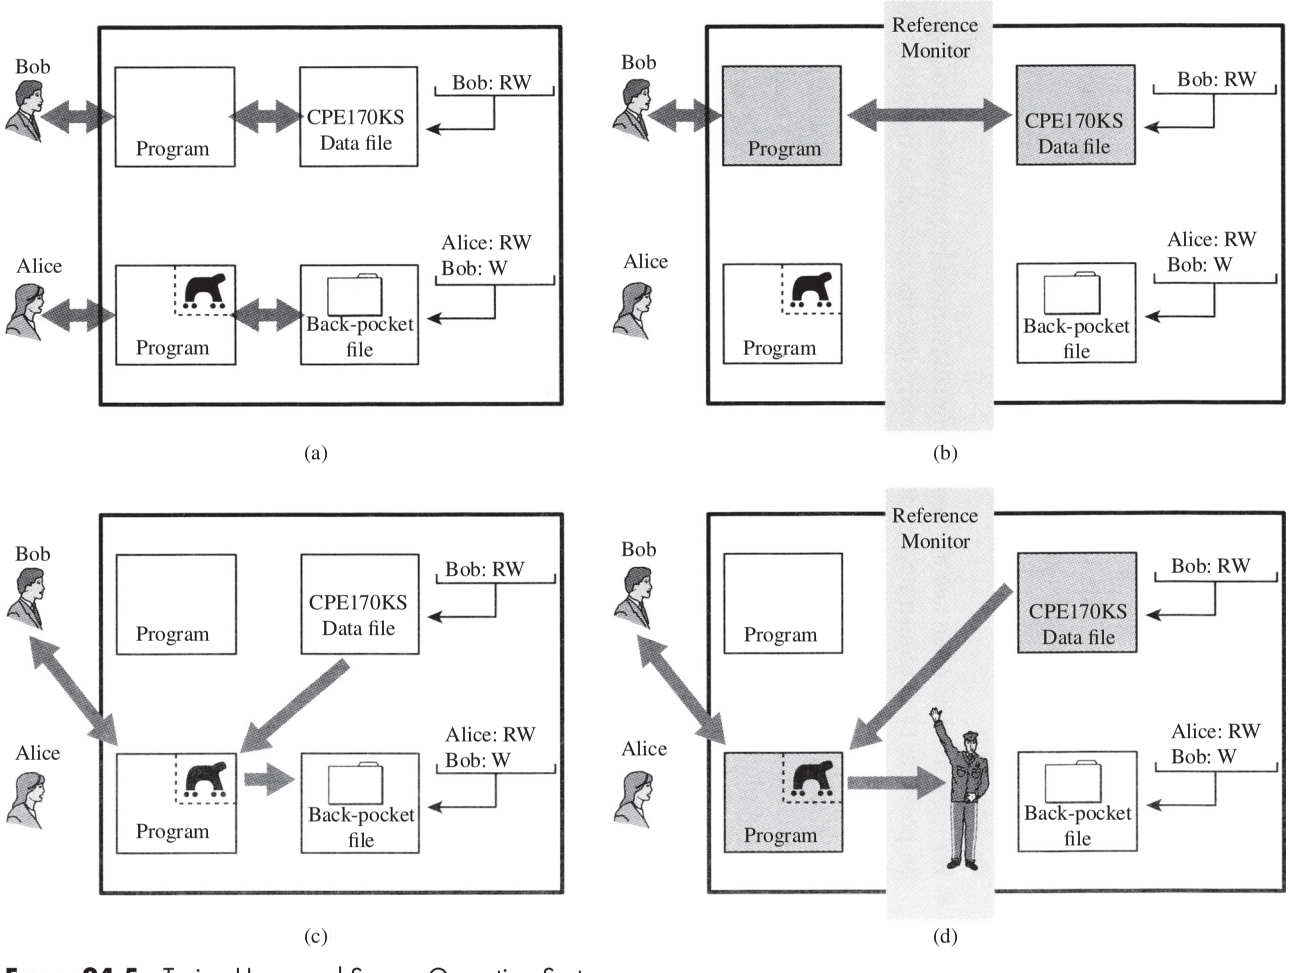


Exhibit 24.5 Trojan Horse and Secure Operating Systems

W2K provides a uniform access-control facility that applies to processes, threads, files, semaphores, windows, and other objects. Access control is governed by two entities: an access token associated with each process and a security descriptor associated with each object for which interprocess access is possible.

24.6.1 Access-Control Scheme. When a user logs on to a W2K system, W2K uses a name/password scheme to authenticate the user. If the logon is accepted, a process is created for the user and an access token is associated with that process object. The access token, whose details are described later, include a security ID (SID), which is the identifier by which this user is known to the system for purposes of security. When the initial user process spawns any additional processes, the new process object inherits the same access token.

The access token serves two purposes:

1. It keeps all necessary security information together to speed access validation. When any process associated with a user attempts access, the security subsystem can make use of the token associated with that process to determine the user's access privileges.
2. It allows each process to modify its security characteristics in limited ways without affecting other processes running on behalf of the user.

The chief significance of the second point has to do with privileges that may be associated with a user. The access token indicates which privileges a user may have. Generally, the token is initialized with each of these privileges in a disabled state. Subsequently, if one of the user's processes needs to perform a privileged operation, the process may enable the appropriate privilege and attempt access. It would be undesirable to keep all of the security information for a user in one systemwide place, because in that case enabling a privilege for one process enables it for all of them.

A security descriptor is associated with each object for which interprocess access is possible. The chief component of the security descriptor is an access-control list that specifies access rights for various users and user groups for this object. When a process attempts to access this object, the SID of the process is matched against the access-control list of the object to determine if access will be allowed.

When an application opens a reference to a securable object, W2K verifies that the object's security descriptor grants the application's user access. If the check succeeds, W2K caches the resulting granted access rights.

An important aspect of W2K security is the concept of impersonation, which simplifies the use of security in a client/server environment. If client and server talk through a remote procedure call (RPC) connection, the server can temporarily assume the identity of the client so that it can evaluate a request for access relative to that client's rights. After the access, the server reverts to its own identity.

24.6.2 Access Token. Exhibit 24.6a shows the general structure of an access token, which includes these parameters:

- **Security ID (SID)**. Identifies a user uniquely across all of the machines on the network. This generally corresponds to a user's logon name.
- **Group SIDs**. A list of the groups to which this user belongs. A group is simply a set of user IDs that are identified as a group for purposes of access control. Each

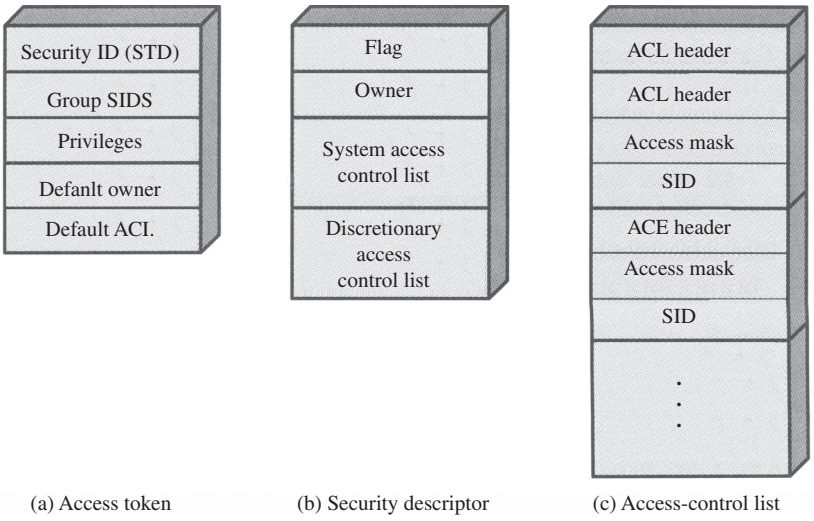


EXHIBIT 24.6 Windows 2000 Security Structures

group has a unique group SID. Access to an object can be defined on the basis of group SIDs, individual SIDs, or a combination.

- **Privileges**. A list of security-sensitive system services that this user may call. An example is create token. Another example is the set backup privilege; users with this privilege are allowed to use a backup tool to back up files that they normally would not be able to read. Most users will have no privileges.
- **Default owner**. If this process creates another object, this field specifies who is the owner of the new object. Generally, the owner of the new process is the same as the owner of the spawning process. However, a user may specify that the default owner of any processes spawned by this process is a group SID to which this user belongs.
- **Default access control list (ACL)**. This is an initial list of protections applied to the objects that the user creates. The user may subsequently alter the ACL for any object that it owns or that one of its groups owns.

24.6.3 Security Descriptors. Exhibit 24.6b shows the general structure of a security descriptor, which includes these parameters:

- **Flags**. Defines the type and contents of a security descriptor. The flags indicate whether the system access-control list (SACL) and discretionary access control list (DACL) are present, whether they were placed on the object by a defaulting mechanism, and whether the pointers in the descriptor use absolute or relative addressing. Relative descriptors are required for objects that are transmitted over a network, such as information transmitted in an RPC.
- **Owner**. The owner of the object generally can perform any action on the security descriptor. The owner can be an individual or a group SID. The owner has the authority to change the contents of the DACL.
- **System access control list (SACL)**. Specifies what kinds of operations on the object should generate audit messages. An application must have the corresponding

privilege in its access token to read or write the SACL of any object. This is to prevent unauthorized applications from reading SACLs (thereby learning what not to do to avoid generating audits) or writing them (to generate many audits to cause an illicit operation to go unnoticed).

- **Discretionary access-control list (DACL)**. Determines which users and groups can access this object for which operations. It consists of a list of access-control entries (ACEs).

When an object is created, the creating process can assign as owner its own SID or any group SID in its access token. The creating process cannot assign an owner that is not in the current access token. Subsequently, any process that has been granted the right to change the owner of an object may do so, but again with the same restriction. The reason for the restriction is to prevent a user from covering his or her tracks after attempting some unauthorized action.

Let us look in more detail at the structure of access-control lists, because these are at the heart of the W2K access-control facility (see Exhibit 24.7). Each list consists of an overall header and a variable number of access-control entries. Each entry specifies an individual or group SID and an access mask that defines the rights to be granted to this SID. When a process attempts to access an object, the object manager in the W2K executive reads the SID and group SIDs from the access token and then scans down the object's DACL. If a match is found—that is, if an ACE is found with a SID that matches one of the SIDs from the access token—then the process has the access rights specified by the access mask in that ACE.

Exhibit 24.7 shows the contents of the access mask. The least significant 16 bits specify access rights that apply to a particular type of object. For example, bit 0 for a file object is File_Read_Data access, and bit 0 for an event object is Event_Query_Status access.

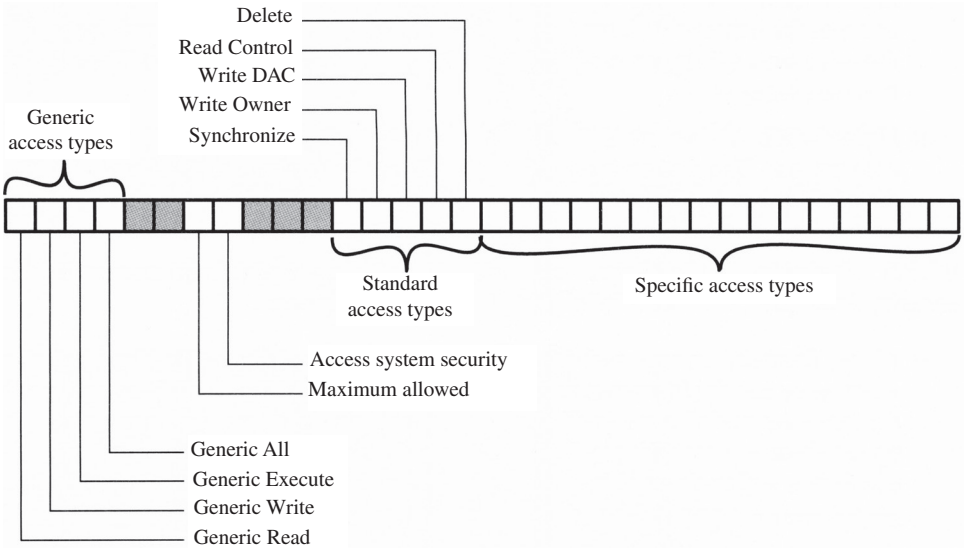


EXHIBIT 24.7 Windows 2000 Access Mask

The most significant 16 bits of the mask contains bits that apply to all types of objects. Five of these are referred to as standard access types:

1. **Synchronize**. Gives permission to synchronize execution with some event associated with this object. In particular, this object can be used in a wait function.
2. **Write_owner**. Allows a program to modify the owner of the object. This is useful because the owner of an object always can change the protection on the object. (The owner may not be denied Write DAC access.)
3. **Write_DAC**. Allows the application to modify the DACL and hence the protection on this object.
4. **Read_control**. Allows the application to query the owner and DACL fields of the security descriptor of this object.
5. **Delete**. Allows the application to delete this object.

The high-order half of the access mask also contains the four generic access types. These bits provide a convenient way to set specific access types in a number of different object types. For example, suppose an application wishes to create several types of objects and ensure that users have read access to the objects, even though read has a somewhat different meaning for each object type. To protect each object of each type without the generic access bits, the application would have to construct a different ACE for each type of object and be careful to pass the correct ACE when creating each object. It is more convenient to create a single ACE that expresses the generic concept *allow read*; simply apply this ACE to each object that is created and have the right thing happen. That is the purpose of the generic access bits, which are:

- **Generic_all**. Allow all access.
- **Generic_execute**. Allow execution if executable.
- **Generic_write**. Allow write access.
- **Generic_read**. Allow read-only access.

The generic bits also affect the standard access types. For example, for a file object, the Generic_Read bit maps to the standard bits Read_Control and Synchronize and to the object-specific bits File_Read_Data, File_Read_Attributes, and File_Read_EA. Placing an ACE on a file object that grants some SID Generic_Read grants those five access rights as if they had been specified individually in the access mask.

The remaining two bits in the access mask have special meanings. The Access_System_Security bit allows modifying audit and alarm control for this object. However, not only must this bit be set in the ACE for a SID, but the access token for the process with that SID must have the corresponding privilege enabled.

Finally, the Maximum_Allowed bit is not really an access bit but a bit that modifies W2K's algorithm for scanning the DACL for this SID. Normally, W2K will scan through the DACL until it reaches an ACE that specifically grants (bit set) or denies (bit not set) the access requested by the requesting process or until it reaches the end of the DACL, in which latter case access is denied. The Maximum_Allowed bit allows the object's owner to define a set of access rights that is the maximum that will be allowed to a given user. With this in mind, suppose that an application does not know all of the

operations that it is going to be asked to perform on an object during a session. There are three options for requesting access:

1. Attempt to open the object for all possible accesses. The disadvantage of this approach is that the access may be denied even though the application may have all of the access rights actually required for this session.
2. Only open the object when a specific access is requested, and open a new handle to the object for each different type of request. This is generally the preferred method because it will not unnecessarily deny access, nor will it allow more access than necessary. However, it imposes additional overhead.
3. Attempt to open the object for as much access as the object will allow this SID. The advantage is that the user will not be artificially denied access, but the application may have more access than it needs. This latter situation may mask bugs in the application.

An important feature of W2K security is that applications can make use of the W2K security framework for user-defined objects. For example, a database server might create its own security descriptors and attach them to portions of a database. In addition to normal read/write access constraints, the server could secure database-specific operations, such as scrolling within a result set or performing a join. It would be the server's responsibility to define the meaning of special rights and perform access checks. But the checks would occur in a standard context, using systemwide user/group accounts and audit logs. The extensible security model should prove useful to implementers of foreign files systems.

24.7 FURTHER READING

Boebert, W., R. Kain, and W. Young. "Secure Computing: The Secure Ada Target Approach." *Scientific Honeyweller* (July 1985). Reprinted in M. Abrams and H. Podell, *Computer and Network Security*. IEEE Computer Society Press, 1987.

Bransted, D. (ed.). *Computer Security and the Data Encryption Standard*. National Bureau of Standards, Special Publication No. 500-27, February 1978.

Denning, P., and R. Brown. "Operating Systems." *Scientific American* 251 (September 1984): 94–106.

Gasser, M. *Building a Secure Computer System*. New York: Van Nostrand Reinhold, 1988.

Gollmann, D. *Computer Security*, 3rd edition. Hoboken, NJ: Wiley & Sons, 2011.

Patterson, D. A., and J. L. Hennessy. *Computer Organization and Design: The Hardware/Software Interface*, 4th ed. Morgan Kaufmann, 2011.

Pfleeger, C. P., and S. L. Pfleeger. *Security in Computing*, 4th ed. Upper Saddle River, NJ: Prentice-Hall, 2006.

Saltzer, J., and M. Schroeder. "The Protection of Information in Computer Systems." *Proceedings of the IEEE* (September 1975).

Sandhu, R., and P. Samarati, "Access Control: Principles and Practice." *IEEE Communications* (September 1994).

Singhal, M., and N. Shivaratri. *Advanced Concepts in Operating Systems*. New York: McGraw-Hill, 1994.

Sinha, P. K. *Distributed Operating Systems: Concepts and Design*. Hoboken, NJ: Wiley-IEEE Press, 1996.

Stallings, W. *Cryptography and Network Security: Principles and Practice*, 5th ed. Upper Saddle River, NJ: Prentice-Hall, 2010.

Stallings, W. *Operating Systems: Internals and Design Principles*, 7th ed. Upper Saddle River, NJ: Prentice-Hall, 2011.

Viega, J., and J. Voas. "The Pros and Cons of Unix and Windows Security Policies." *IT Professional* 2, no. 5 (September/October 2000): 40–47.

24.8 NOTES

1. P. Denning and R. Brown, "Operating Systems," *Scientific American* (September 1984).
2. J. Saltzer and M. Schroeder, "The Protection of Information in Computer Systems," *Proceedings of the IEEE* (September 1975).
3. C. P. Pfleeger and S. L. Pfleeger, *Security in Computing*, 4th ed. (Prentice-Hall PTR, 2006).

CHAPTER 25

LOCAL AREA NETWORKS

N. Todd Pritsky, Joseph R. Bumblis, and Gary C. Kessler

25.1 INTRODUCTION 25·1

25.2 POLICY AND PROCEDURE ISSUES 25·1

25.3 PHYSICAL SITE SECURITY 25·3

25.4 PHYSICAL LAYER ISSUES 25·4

- 25.4.1 Sniffers and Broadcast LANs 25·4
- 25.4.2 Attacks on the Physical Plant 25·5
- 25.4.3 Modems, Dial-Up Servers, and Telephone Lines 25·5
- 25.4.4 Wireless LAN Issues 25·7

25.5 NETWORK OPERATING SYSTEM ISSUES 25·9

- 25.5.1 Windows 9x 25·10
- 25.5.2 Windows NT/2000, XP, Vista, Win7, and Win8 25·11
- 25.5.3 UNIX 25·13
- 25.5.4 MacOS 25·15

25.6 CONCLUSION 25·16

25.7 FURTHER READING 25·16

25.8 NOTES 25·17

25.1 INTRODUCTION. This chapter discusses generic issues surrounding local area network (LAN) security. Securing the LAN is essential to securing the Internet because LANs are where most of the attackers, victims, clients, servers, firewalls, routers, and other devices reside. Compromised LAN systems on the Internet open other nodes on that local network to attack and put other systems at risk on the Internet as a whole. Many of the general issues mentioned herein are described in more specific terms in other chapters of this *Handbook,* such as Chapters 15, 22, 23, and 47 in particular.

25.2 POLICY AND PROCEDURE ISSUES. Thirty years ago, nearly all computer users had accounts on a shared mainframe or minicomputer. A single system manager was responsible for security, backup, disaster recovery, account management, policies, and all other related issues. Today all users are system managers, and, in many cases, individuals have responsibility for several systems. Since the vulnerability of a single computer can compromise the entire LAN, it is imperative that there be rules in place so that everyone can work together for mutual efficiency and defense. But where policies and procedures can be centralized, they should be, because most users do not take the security procedures seriously enough.

The next list, modified from the Internet Engineering Task Force (IETF) Request for Comment (RFC) 2196, is a rough outline of LAN-related security policies and procedures that should at least be considered.[1]

1. Administrative Policies Framework
   a. Information security issues
      i. Password management procedures
         1. Are passwords assigned or chosen by user?
         2. Are password auditing procedures in place?
         3. Are password policies enforced (e.g., minimum length, allowed and required characters, expiration, blacklisting)?
         4. How many passwords are required to access all systems?
      ii. Virus protection
         1. Are servers protected?
         2. Is there an e-mail "viruswall"?
         3. Is virus protection a centrally managed activity, or up to each user?
         4. How do users maintain the current virus signature database?
      iii. Encryption and certificates
      iv. Security event handling
   b. Network connectivity issues
      i. Dial-up access
      ii. Hubs versus switches
      iii. Identification and authentication
         1. Passwords and management
         2. Authentication systems
         3. Two-factor authentication?
         4. Single sign-on?
         5. Biometrics
   c. Physical site security and disaster recovery
      i. Physical security of systems
      ii. Dial-up access lines and modems
      iii. Scheduling backups
      iv. Access to servers
      v. Storing and limiting access to backup media
      vi. Disaster recovery and contingency plans
   d. Operating system and LAN security
      i. Operating system-specific issues
         1. Monitoring operating system security vulnerabilities
         2. Applying security patches
         3. Securing the operating system (OS) during installation
         4. Auditing the systems

ii. Is Dynamic Host Configuration Protocol (DHCP) employed?
iii. Log analysis software and procedures
iv. Vulnerability testing
v. Intrusion detection tools

2. User Policies Framework
   a. Written, published network security policy
   b. Network/Internet appropriate use policy (AUP)
   c. User training and education
      i. Security (general issues)
      ii. Importance of protecting customer data, client information, and other private information such as medical claim records and patient information
      iii. Policies and AUPs
      iv. Suggestions for how to be safe
   d. Virus protection
      i. Using antivirus software
      ii. Maintaining a current virus signature database
   e. Choosing good passwords
   f. Best practices
      i. Email (netiquette and handling attachments)
      ii. Browser (ensure that Java, JavaScript, ActiveX, and other auto-execution code is current as specified by the manufacturer. See "Notes"[2] for helpful links.)
      iii. Microsoft Office Suite
         1. Use of macros
         2. Implications for document management
      iv. Protecting files on the server and your own system
         1. Use of Windows "shares" and UNIX Network File System (NFS)
         2. Use of NetWare, New Technology File System (NTFS), and UNIX access controls
            i. Spotting a possible compromise (methods for identifying intrusions or other unauthorized activity)
            ii. What to do and whom to contact if a compromise is suspected

Not every issue will apply to all networks, but each LAN is an ever-evolving entity, and the policies guiding its operation also must evolve.

25.3 PHYSICAL SITE SECURITY. The physical protection of the site is very important but too often overlooked. One reason that site security is often lacking is that some of the policies and procedures can be perceived as messages to employees that they are not trusted. Nevertheless, physical site security includes many aspects of protecting network servers, communications facilities, individual user's systems, and information.

For more information on facilities security, see Chapters 22 and 23 of this *Handbook.*

25.4 PHYSICAL LAYER ISSUES. The LAN itself has a number of additional vulnerabilities because the systems and media are so widely dispersed. This section discusses securing the LAN infrastructure.

25.4.1 Sniffers and Broadcast LANs. Traditional LAN media access control (MAC) schemes operate assuming a logical, if not physical, broadcast topology. In a broadcast network, every station hears every transmission. When operating in *promiscuous mode*, a LAN station will read every frame that goes by regardless of whether the frame is addressed to the station or not.

When protocol analysis, or *sniffer*, software became available in the 1980s, every station on the LAN became a potential network analysis and management tool—or a surreptitious eavesdropper. Since so many of the network's applications—particularly Transmission Control Protocol/Internet Protocol (TCP/IP)-based applications—transmit passwords and files in plaintext, this type of software is potentially very dangerous.

A number of sources offer very powerful and very flexible commercial packet-sniffing software, such as Network Associates' SnifferPro and Novell's LANalyzer. These packages usually have additional capabilities, such as network monitoring, performance monitoring, and traffic analysis. Prior to 1990, network protocol analysis required a special piece of hardware with a connection to the network. Today, a large number of software packages that do TCP/IP packet sniffing can help an intruder because they can be installed directly on an individual's laptop or desktop computer. Some of these packages include:

- WireShark (Windows, MacOS, and UNIX)
- BUTTsniffer (Windows NT)
- Network Monitor (free with, and for, Windows NT)
- Sniffit (Linux, SunOS, Solaris, FreeBSD, Irix)
- snort (UNIX and Windows)
- Solsniff (Solaris)
- tcpdump (MacOS and UNIX)
- WinDump (DOS)

Relatively few countermeasures can be taken against these kinds of tools. Fortunately, they are effective only on a broadcast network, such as a hubbed LAN. If an Ethernet hub, for example, is replaced with a switch, the only traffic broadcast to all hosts on the network are those frames actually addressed to the LAN's broadcast address. In this case, a station can sniff only that traffic that is going to or coming from the station with the sniffer software. Replacing all hubs with switches may be unreasonable in many environments, but placing all servers on a switch instead of a hub will improve both performance and security.

Other network-based tools that can detect a host with a network interface card (NIC) in promiscuous mode, such as AntiSniff (Windows) and sentinel (UNIX). These tools work by performing a number of tests to detect the promiscuous host; they then check the network host's operating systems, domain name system (DNS) activity, and network and machine latency.

Sniffers can be defeated using cryptography. Use of secure IP (IPsec) and Secure Shell (SSH) for TCP/IP applications can provide privacy and integrity to all

communication and applications. Information on IPsec is available on www.ipsec-howto.org/x202.html, and SSH is available from SSH Communications Security (www.ssh.com).

25.4.2 Attacks on the Physical Plant. The most common medium employed today on LANs is copper-based, unshielded twisted pair (UTP) cable. All copper media, including UTP and coaxial cable, emanate a magnetic field because of the changing current on the wire. Van Eck monitoring devices can pick up these emanations remotely and reproduce the frames on the wire or keystrokes at a machine. As far-fetched as this might sound, the vulnerability is very real. The U.S. Government and military have a set of standards to reduce and limit electromagnetic radiation (EMR) called TEMPEST, and the National Security Agency (NSA) has a list of products that are TEMPEST-compliant. Indeed, this vulnerability may not be a major problem in most organizational networks, but there is some set of networks where this is a concern.

An alternative to encasing workstations in Faraday cages (i.e., copper mesh layers surrounding all components) is to generate random electronic noise that masks meaningful radiated data transmissions.[3]

One way to reduce or eliminate EMR is to reduce or eliminate the amount of copper media on the network. Optical fiber, for example, has no EMR because there is no electricity on the line. It does not eliminate the EMR at the desktop itself, but it does prevent its entry into all interconnecting fiber cables.

Users also can be the source of some types of denial-of-service (DoS) attacks, either purposely or accidentally. Consider coaxial cable Ethernet networks (10BASE-5 or 10BASE-2) where nodes are attached to a common LAN medium. In both cases, the coaxial cable has a terminating resistor at the end of the wire to eliminate reflection of the signal. If the resistor is removed, it allows extra noise on the wire that can block all traffic on the network. End-of-cable resistors should be beyond the reach of users if at all possible.

Similar DoS attacks may occur whenever users modify the wiring scheme of the network. Removal of terminating resistors is only one action that can cause problems. Hubs in common areas might be unplugged or have network connectors removed. A token ring hub-to-hub connection can be detached, breaking the integrity of the ring and thus preventing LAN hosts from communicating with other hosts and denying service to users.

It is important that the physical network should be secured to the greatest extent possible. LAN managers should educate users to avoid the accidental problems that might occur and even how to recognize nefarious attacks.

25.4.3 Modems, Dial-Up Servers, and Telephone Lines. Modems anywhere on the LAN are a potential danger, particularly those that are connected directly to a user's system, with or without official sanction. Modems can provide a back door into the LAN, possibly bypassing the firewall, strong authentication, proxy server, and any other network security. Although less common in today's LAN architectures, dial-up modems are still used to dial into routers/bridges during network troubleshooting exercises. Care, as discussed in this section, should still be taken if dial-up modems are in use to access LAN devices.

In general, all modems should be concentrated at the network's dial-up server. Individual user systems should, in most cases, be banned from having modems on the network. This is a difficult rule to enforce, however, because laptops and an increasing number of desktop systems include preinstalled modems, and a user can always connect

an external modem to the serial port of any system. This is an example of why security managers have to integrate policies into the culture of an organization. Otherwise, users will find a way around what they perceive to be prohibitive and onerous policies, and modems are one way to circumvent the corporate firewall.

Modems in auto-answer mode are particularly dangerous. Although most companies do not advertise their dial-in telephone numbers, these numbers do not remain secret for very long from someone who wants to find them. Anyone with a telephone directory can easily start mapping an organization's block of corporate telephone numbers. For example, if the main number is 802-555-3700, attackers have a place to start. When attackers call the main number and ask the receptionist for the organization's fax number, they obtain even more information. Using war-dialer software, attackers can scan an entire block of telephone numbers (e.g., 555–3700 through 555–3799) and obtain a list of which numbers are active, which respond with a tone, and what the tone represents (e.g., fax, modem, etc.). If a user has an auto-answer modem on a computer, attackers may gain access to the user's system without so much as a password. Security managers should work with their local telephone companies to obtain telephone numbers for modem lines that are not in the organization's telephone block.

The dial-up server, then, is the place to concentrate the modems and the authentication of dial-up users. There are several strategies for authentication at the dial-up server, the strongest being some form of two-factor authentication, such as a combination of passwords and a token. Another strong protection mechanism is to implement a dial-back mechanism, so that once a bona fide user logs in, the system hangs up and calls back to a preconfigured telephone number. This is a very effective scheme that works well with fixed-location telecommuters but not with roaming employees. In addition, attackers have been known to tamper with the central switch of a telephone company to call-forward from an assigned number to the attacker's own modem.

When a user requires two separate logons, one to the dial-up server and then one to a domain server, security restrictions should control what the caller is allowed to do after passing the first test. One of the authors of this chapter worked with a company that had a shared secret (an oxymoron) telephone number for its modem bank and then a single shared username and password for all the users to authenticate to the dial-up server. To access files and shared network resources, the user then had to authenticate to the domain controller. But after passing the identification and authentication for the first server, an attacker was on the organization's LAN and had complete, unfettered access to the Internet and an identity that implicated the company in any possible malfeasance.

Some guidelines for securing dial-up servers include:

- Maintain and monitor telephone logs.
- Configure all software and modems so that a user logout forces the modem to disconnect, and a modem disconnect forces a user logout.
- Configure the modems so that they return to their default configuration after every connection.
- Implement a dial-back mechanism, where possible.
- Use two-factor authentication for roaming users.
- Periodically scan internal telephone numbers for unauthorized modems.
- Prevent the display of any banner message when the user connects to the modem, and certainly do not display any sort of *welcome* message.

- Train the organization help desk in social engineering techniques, and prohibit them from giving out modem telephone numbers, user names, or other sensitive information that could help attackers.

25.4.4 Wireless LAN Issues. Wireless LANs (WLANs) have vulnerabilities their wired counterparts do not. The most obvious difference between wired and wireless networks is the medium itself. Although copper-based LANs emit a small amount of radiation that can be intercepted, the entire basis of wireless LANs is transmitting data using relatively strong radiation in some form.

There are WLANs based on infrared signals that cannot penetrate building walls and so achieve some degree of security due to the limited propagation of those signals. Such LANs typically are found in networks requiring high levels of security. However, most WLANs today use radio transmission techniques. In these networks, anyone on a nearby street can use a listening device to intercept data and even capture the network identifiers required to connect to the LAN. The practical range of interception is governed by the inverse-square law for the signal strength (it declines as the square of the distance from the source) and by the sensitivity and signal-to-noise characteristics of the receivers.

Fortunately, there is a certain measure of security within the physical layer itself. The Institute of Electrical and Electronics Engineers (IEEE) 802.11-based LANs employ either direct-sequence spread spectrum (DSSS) or frequency-hopping spread spectrum (FHSS) techniques. As always, depending on range and noise levels, it is possible to eavesdrop. However, interpreting the signals is made more difficult by how DSSS and FHSS work.

To make sense of the transmissions, the receiver must know either the *chipping code* used in a DSSS network or the *frequency-hopping pattern* in an FHSS implementation. Without such information, the signal will appear to be nothing more than background noise in the industrial, scientific, and medical (ISM) radio bands to the illicit receiver. It is not an insurmountable problem for the would-be eavesdropper, but the work factor is greatly increased when compared to narrowband radio techniques. The spread spectrum approach also offers more reliability in the face of interference from denial of service (i.e., intentional jamming), as the signal is spread over a broad range of frequencies. Some vendor equipment also comes with software components that allow for tuning around interference.

For example, the IEEE 802.11g amendment uses the newer orthogonal frequency division multiplexing (OFDM) from 802.11a for higher data speeds, yet is backward compatible with 802.11b using DSSS, which was already using the same ISM frequency band. DSSS data rates of 1, 2, 5.5, and 11 Mbps are supported, as are OFDM data rates of 6, 9, 12, 18, 24, 48, and 54 Mbps. IEEE requires only mandatory data rates of OFDM using 6, 12, and 24 Mbps, regardless whether it is 802.11a or 802.11g OFDM.[4]

For greater privacy than that provided by the physical layer alone, the 802.11 standard includes an optional encryption method called Wired Equivalent Privacy (WEP). This technique is truly optional so not all vendors support the standard. WEP uses a 40-bit form of the RC4 algorithm by default, although some products support stronger, 128-bit versions. It is a good idea to choose a product that offers more than just the 40-bit version, as a 40-bit keyspace does not provide great security given today's computing power.

Because WEP does not offer strong encryption and does not describe a standard key-exchange mechanism, many vendors have implemented layer three tunneling methods, such as those found in virtual private networks (VPNs), to provide greater privacy. These

VPN-based approaches generally employ other encryption processes (e.g., Microsoft Point-to-Point Encryption) as used in the Point-to-Point Tunneling Protocol (PPTP) that use longer keys than WEP and often support Public Key Infrastructure (PKI) or other key-exchange mechanisms. Some implementations also provide authentication through standards such as Remote Access Dial-In User Service (RADIUS) for more flexible client management. The major problem is that these approaches are not all interoperable and are not necessarily multiprotocol capable.

WEP also can be used for authentication to prevent unauthorized access to the WLAN itself. Such authentication adds another layer of protection to the username and password combination employed by typical server software. Before gaining access to information resources on the server, a client first must gain access to the physical medium. Using the shared key scheme, a wireless device must possess the same encryption key as the LAN's access point, the device enabling wireless connectivity to the wired portion of the LAN. Any data transmitted must be encrypted with the key, or the frame will be ignored. Many wireless access products also have the capability to create access control lists based on MAC addresses to filter LAN connections.

IEEE 802.11i-2004 or 802.11i security standard, implemented as Wi-Fi Protected Access II (WPA2), is an amendment to the original IEEE 802.11. WPA2 was originally developed by the Wi-Fi Alliance (www.wi-fi.org). This standard specifies security mechanisms for wireless networks. It replaced the short Authentication and Privacy clause of the original IEEE 802.11 standard with a detailed Security clause. In the process it deprecated the original WEP encryption method. The amendment was later incorporated into the published IEEE 802.11-2007 standard.

In summary, network managers should consider these security items when evaluating wireless network components:

Physical Layer Schemes

- **Infrared**. Cannot penetrate walls and is good for high-security applications.
- **FHSS**. Signal hopping provides good level of security but complexity of technique limits bandwidth.
- **DSSS**. Low "spreading ratio" can increase available bandwidth but also the possibility of interception and jamming.
- **OFDM**. Used for higher data rates in newer WLANs but may be more sensitive to noise and jamming.

Encryption Options

- **WEP**. Deprecated in 2004, commonly employed 40-, 64-, or 128-bit keys with the Rivest Cipher 4 (RC4) algorithm.
- **Wi-Fi Protected Access II**. A more secure alternative to WEP, employing the Advanced Encryption Standard (AES) Counter-Mode/CBC-MAC Protocol (CCMP).[5]
- **Alternative encryption methods.** Often more secure than WEP and WPA2 but not always interoperable.

Authentication Methods

- **Wired Equivalent Privacy.** Uses pre-shared key (PSK) mode, which requires clients to have the same encryption key as the LAN access point, which introduces key management issues.
- **Wi-Fi Protected Access II.** WPA2 supports IEEE 802.1X/EAP authentication or PSK technology.

- **Access control list.** Allows only certain clients to gain physical access to the LAN, based on the MAC address; this adds complexity to client management.
- **Server-based authentication.** Flexible user authentication with RADIUS for an additional layer of protection from illicit connections.

For more information about security for wireless networks, see Chapter 33 in this *Handbook*.

25.5 NETWORK OPERATING SYSTEM ISSUES. In the early 1990s, it was common to find desktop systems running the Windows operating system *and* the Novell NetWare network operating system (NOS). Desktop applications ran over Windows, and NetWare was used only to move files to and from the shared file space, or to print documents.

Today, the distinction among the desktop operating system, server operating system, and NOS has disappeared. Operating systems such as Linux, MacOS, UNIX, and Windows all provide desktop application suites with networking capabilities, including communications protocols such as TCP/IP. There are some general security considerations for all LANs regardless of the specific operating system:

- Use whatever capabilities are provided by the operating system to employ strong passwords.
- Create password policies that force strong passwords. Change passwords periodically and do not allow reuse. Periodically audit passwords using password-cracking tools such as L0phtCrack (Windows) or crack (UNIX). Ensure that the administrator and root accounts are given passwords that are not widely distributed or guessed.
- Disable (or uninstall) any services that are not being used.
- Keep the operating system and application software up to date, and with the latest security patches installed.
- Carefully manage access control lists for files and other system/network resources.
- Strictly define users, groups, and network/domain trusts.
- Tightly secure any running applications.
- Log on as administrator or root only when necessary; otherwise log on as a regular user.
- Allow operators and administrators to log on only locally at server systems.
- Limit use of guest, demo, or anonymous accounts.
- Where feasible, put boot and system files as well as application files and data on different partitions, hard drives, or input/output (I/O) controllers.
- Regularly audit server systems.
- Monitor log files.
- Remove floppy, CD, DVD, and thumb drives from servers after the system has a stable configuration.
- Implement best industry practices when securing the operating system.
- Use vulnerability assessment tools on a regular basis to scan servers.
- Use intrusion detection tools to monitor potential attacks on the LAN that are launched from the internal network.

- If using the Simple Network Management Protocol (SNMP) for network administration, carefully choose community names and block external access to the SNMP service. Make management information bases (MIBs) read-only, where possible.
- Avoid use of the DNS HINFO (host info) resource record to identify the central processing unit type and the installed operating system.

Specific operating system vulnerabilities are beyond the scope of this chapter; entire books and Websites are devoted to securing some of these individual operating systems. At a minimum, network managers must monitor their operating system vendor's Website and all other sites that cover the NOS's security. The sections that follow provide some general observations and comments about the various network operating systems.

25.5.1 Windows 9x. All of the Windows operating systems (including NT and 2000) support peer-to-peer resource sharing and are vulnerable to exploitation of NetBIOS file and print sharing. In particular, when file and print sharing is enabled, the service is, by default, bound to TCP/IP. Although this does not cause an additional exposure to systems on the local network (since shares can be seen by other nodes on the LAN anyway), it does provide a potential vulnerability for hosts connected to the Internet. File and print sharing can be unbound from TCP/IP using Start, Control Panel, and Network.

Windows 9x (including Windows 95, Windows 98, and Windows ME) systems also have a vulnerability in the way authentication is performed when a user wishes to access a remote share. Windows uses the challenge-handshake authentication protocol (CHAP) for all passwords that need to be sent over the network, so passwords never appear in plaintext on the LAN. However, because Windows 9x uses the *same* challenge during a given 15-minute period, an intruder with a physical access to the LAN and to a sniffer could effect a replay attack by resending a duplicate authentication request and remapping a share on the Windows 9x system. This example illustrates the critical role of physical security in preventing compromise of LANs.

One of the classic Trojan horse programs for Windows 9x is Back Orifice (BO). Advertised as a remote Win9x administrator tool, it can be used by a nefarious user to take total control of someone else's system, including the capability to modify the registry, reboot the system, transfer files, view cached passwords, spawn processes, and create shares. SubSeven and NetBus are other tools that can be used to take control of a remote Windows system. Some commercial virus scanners can detect these programs on individual systems, and several workstation firewalls exclude communications characteristic of these Trojans.

Windows 9x has no particular logon security mechanism. Although every user might be forced to log on to a system, any user can log in with any username and any password to get at least basic access. Several password-cracking programs are available through the Internet to break Windows' .PWL password files, which are accessible once a user has access to the network. If a password-protected screen saver is preventing an attacker from logging in, there is an easy way around this as well: Simply reboot the computer. However, third-party security programs include nonbreakable secure logins and secure screen savers; many include bootlock functions that prevent any access whatever unless a valid password is entered. Current examples of such software can be located easily using buyers' guides such as the annual listing from the Computer Security Institute.[6]

25.5.2 Windows NT/2000, XP, Vista, Win7, and Win8. From a security perspective, Windows 2000 Millennium Edition (ME) is not significantly stronger than Windows 9x. Network administrators should periodically scan the network's public shares to ensure that they are appropriate. Windows NT Server, NT Workstation, and 2000 Server editions are built for security, and network services and have the software architecture to support these services. However, a security vulnerability announcement related to these operating systems seems to come out almost weekly. Many of the hacking tools available for Win9x also are available for Windows NT and 2000; Back Orifice 2000 (BO2K), for example, is an NT/2000 version of BO and NetBus also can take control of an NT/2000 host.

Scripting holes in Internet Explorer (IE) and Office 2000 make all Windows systems susceptible to a variety of new virus and worm attacks. Although the early viruses, such as Melissa and I LOVE YOU, required users to open e-mail attachments, that is no longer so. Microsoft Outlook and Outlook Express will execute Hypertext Markup Language (HTML) and script code in the body of an e-mail by default. Several ActiveX components also will execute from an e-mail containing HTML and script code; examples of such controls include Scriplet.typlib (ships with IE 4.x and 5.x) and the UA control (Office 2000). The best protection against these types of vulnerabilities is to define Outlook and Outlook Express to read e-mail in the "Restricted Sites Zone" and disable all Active Scripting and ActiveX related settings in that zone. This vulnerability affects all Windows systems with Office 2000 or Internet Explorer 4.x/5.x installed, even if IE is not used.

Securing Windows NT/2000 systems is well beyond the scope of this chapter, but some of the precautions, in addition to those listed already, follow.

- Format the drive using NTFS rather than FAT.
- Use long file names and disable the DOS 8.3 naming format.
- Disable the Everyone group.
- Rename the administrator account.
- Turn auditing on. (It is off by default.)

All NT-based systems have been given the C2 security rating by the National Computer Security Center, which includes Windows NT 3.5 and 4.0 and Windows 2000 SQL Server version 8.0. This means that when these versions of Windows are installed correctly, they meet evaluation criteria set forth in the National Computer Security Center (NCSC) *Orange Book* of security specifications. C2 certification is not applied to the operating system itself; rather, it is applied to a particular installation. Microsoft provides tools to audit a site so administrator(s) can deploy the correct hardware and software configurations to achieve this level of security in a network.

Windows 2000 introduced a new feature that administrators might want to employ. Windows NT introduced the ability to compress and decompress files on the fly. Windows 2000 introduces the capability to encrypt and decrypt files on the fly. The Encrypting File System (EFS) uses public key cryptography and the Data Encryption Standard (DES) to encrypt files on the hard drive. EFS includes an optional capability for a recovery mechanism in case of key loss. Organizations using EFS on critical systems should consider employing this mechanism to protect against loss or destruction of the private key.

Windows XP, which has sold well over 400 million copies and still enjoys the second largest installed base of any Microsoft OS next to Win7, introduced a number of

security-oriented features. The major improvement was the inclusion of a firewall capability in both the Home Edition and Professional versions. Professional goes a couple of steps further by offering EFS file encryption, Kerberos,[7] smart card support, and a new software restriction feature allowing administrators to mitigate the impact of viruses and Trojan horses. XP also supports raw sockets, which in itself is not unusual—UNIX and Linux do as well. This feature is intended to increase the functionality of Internet services, but in Microsoft's implementation, it is available to any user, no matter what privilege level. Thus hackers can possibly gain control of a computer running XP and use it to initiate DoS attacks by commanding the OS to generate a flood of traffic.

Although Microsoft's Windows Vista was released with much fanfare—and after a number of delays—it has not been as widely implemented as might have been due to concerns about security and stability of the platform, not to mention a perception that many of the security features create a great deal of inconvenience in the user experience (e.g., User Account Control, to be described, and the Digital Rights Management implementation). Regardless, a prime mover of Vista development was Microsoft's Trustworthy Computing initiative, and as such, a number of new security capabilities were added.

Features include User Account Control (UAC), Bitlocker Drive Encryption, Windows Defender, Data Execution Prevention, Application isolation, Windows Service Hardening, Network Access Protection, and a variety of others. UAC is the most visible from the user perspective, requiring user intervention before allowing any action that requires administrative-level privileges. These include[8]:

- Right-clicking an application's icon and clicking "Run as administrator"
- Changes to files or folders in %SystemRoot% or %ProgramFiles%
- Installing and uninstalling applications
- Installing device drivers
- Installing ActiveX controls
- Changing settings for Windows Firewall
- Changing UAC settings
- Configuring Windows Update
- Adding or removing user accounts
- Changing a user's account type
- Configuring Parental Controls
- Running Task Scheduler
- Restoring backed-up system files
- Viewing or changing another user's folders and files

As an interesting historical aside, the National Security Agency assisted Microsoft in the development of Vista. As reported in the *Washington Post* in January 2007:

> For the first time, the giant software maker is acknowledging the help of the secretive agency, better known for eavesdropping on foreign officials and, more recently, U.S. citizens as part of the Bush administration's effort to combat terrorism. The agency said it has helped in the development of the security of Microsoft's new operating system—the brains of a computer—to protect it from worms, Trojan horses and other insidious computer attackers. …
>
> The NSA declined to comment on its security work with other software firms, but Sager said Microsoft is the only one "with this kind of relationship at this point where there's an acknowledgment publicly."

> The NSA, which provided its service free, said it was Microsoft's idea to acknowledge the spy agency's role. …
>
> "I kind of call it a Good Housekeeping seal" of approval, said Michael Cherry, a former Windows program manager who now analyzes the product for Directions on Microsoft, a firm that tracks the software maker.
>
> Cherry says the NSA's involvement can help counter the perception that Windows is not entirely secure and help create a perception that Microsoft has solved the security problems that have plagued it in the past. "Microsoft also wants to make the case that [the new Windows] is more secure than its earlier versions," he said.[9]

It is left as an exercise for the reader to decide whether having a spy agency working on a premier OS is a good thing or not.

One of the most important capabilities for the network security manager is to audit the Windows server systems to protect their integrity. These tools are part of the base operating system or the Windows NT Resource Kit:

- *netstat* examines open ports.
- *Event Viewer* examines application, security, and system logs.
- *net start, net user, net group, net local group* display running services, users, groups, and local groups.
- *dumpel* converts Event Viewer logs to simple text files.
- *NetMon* displays network traffic.
- *netsvc* displays local and remote running services and drivers.
- *addusers* displays users and groups.
- *findgrp* displays local and domain groups for a user.
- *local* and *global* show all members of specific local or global groups.
- *dommon* displays trusted domains.
- *xcacls* examines the file Access Control Lists (ACL).
- *perms* examines the ACLs associated with a user.
- *sysdiff* displays changes in the Registry and file system.
- *regdmp* creates an ASCII version of the Registry.
- *ralist* lists a domain's Remote Access Servers (RAS).
- *rasusers* lists users authorized for dial-in access.

According to Tech Review Source (www.techreviewsource.com /Windows-7/windows-7-security-features#.UL-S3WdN_f0), Windows 7 was the most secure OS Microsoft had released to date. Security features include the new Action Center, improved UAC, Parental Controls, Internet Explorer 8 Security, Windows Defender, and Microsoft Security Essentials. Windows 8 added many changes to the operating system but they were primarily in the user interface and user experience; Win8 capitalized on Win7's newer security capabilities.

25.5.3 UNIX. UNIX is the oldest operating system still in widespread (and growing) use today. Readers needing guidance for Novell Netware should refer to Chapter 18 in the fourth edition of this *Handbook*. Originally developed in 1969 at AT&T Bell Laboratories, UNIX became the first operating system to integrate network communications when TCP/IP was bundled into Berkeley Software Development

(BSD) 4.2 UNIX in 1984. UNIX had traditionally been reserved for server systems and hardcore computer users. With the development of the X-Windows interface for UNIX and the wide deployment of Linux since the mid-1990s, UNIX and its variants represent the only significant competition to Windows in the desktop and server environment.

Like TCP/IP and the Internet itself, UNIX was developed for functionality and use within a trusted user community. As such, while UNIX has many powerful tools, it does not have a cohesive security architecture, nor is it an inherently secure operating system.

UNIX has most of the basic operating system protections: passwords, access control lists, groups, user privilege levels, and so on. But UNIX also comes with a large variety of services (daemons) enabled by default, including *File Transfer Protocol (FTP), Telnet, finger, echo, chargen, daytime, Remote Procedure Call (RPC), BIND*, and more. In addition, nearly every UNIX daemon has had some sort of security vulnerability reported at one time or another, with buffer overflows being quite prevalent.

There are many things that an administrator should consider when securing a UNIX/Linux system. In addition to the general steps just listed, the security manager might also:

- Disable (or remove) any unused services, particularly *finger*, the BIND name daemon (*named*), *RPC, sendmail,* Trivial FTP (tftp), Post Office Protocol (POP), Internet Message Access Protocol (IMAP), *sadmind, mountd*, and Network File System (NFS).
- Install the latest version and security patch of all installed software.
- Take great care when configuring access control lists and other sharing.
- Prevent running *Sendmail* in daemon mode (turn off the -bd switch) on machines that are neither mail servers nor mail relays.
- Limit use of the "r" remote access protocols.
- Use shadow password files.
- Implement TCP Wrappers to control access to services.
- Consider using encrypted communication protocols, such as Secure Shell (SSH) or Secure Sockets Layer (SSL), for remote access. Prevent transmission of cleartext passwords over the Internet.

One of the most important capabilities for the network/security manager is to audit the Windows server systems to protect their integrity. These tools are part of the base operating system or the Windows NT Resource Kit:

- *netstat* examines open ports.
- *lsof* displays hidden file space and network connections.
- *tcpdump* displays network traffic.
- *who* displays users that are logged on and the *utmp* log file.
- *last* displays login history, and the *wtmp* log file.
- *lastb* display a history of bad logins (and the *btmp* log file).
- *syslogd* is a central server facility for managing and logging system messages.
- *TCPWrapper* monitors and manages incoming service requests.

25.5.4 MacOS. The Macintosh operating systems have mostly given way to other platforms. That said, they are still worth mentioning, as Apple's market share across the board has been buoyed somewhat in recent years by its strength in mobile and tablet devices. The Macintosh was the first desktop operating system that included networking and resource sharing via AppleTalk as integral elements. But like UNIX and TCP/IP before it, the MacOS is designed for convenience and usability but not for security.

Because of its peer-to-peer nature, MacOS has traditional had a number of potential exposures. Although Windows and UNIX also could operate in a peer-to-peer mode, a novice user generally would not know how to share resources and thereby might not inadvertently open holes.

For example, a nefarious Mac network user could quickly see what servers and shares were available on the network by using the original *Chooser* accessory. In desktop use, Macs had relatively little security. Password protection was provided by default only with some laptop systems, and not even a password-protected screen saver came standard with the system. In short, there was very little standing between a determined attacker and a Mac computer. Third-party software was required to provide password protection against access to the system and files, or for data protection with disk encryption.

Mac-based viruses and worms continue to be much less prevalent than their Windows counterparts, but Macs are not totally immune. First, those viruses that depend on Microsoft Office Suite software will work because the Mac versions of Word and Excel employ macros. Second, Internet-based attacks aimed at TCP/IP—such as the Ping-of-Death, Teardrop, and SMURF—can still affect a Mac server. There are more choices for add-on security than in the past, with packages like Norton™ Antivirus for Mac® and Norton™ Internet Security for Mac® being similar to their counterparts for other operating systems.

The Mac OS has undergone a fundamental transformation since the last "classic" release of OS 9 back in 1999. With the advent of OS X, Apple introduced its first UNIX-based platform, which has been preloaded on all Macintosh computers since 2002.

Version 10.5 (Leopard) was released in 2007 and added a number of security enhancements to address traditional platform weaknesses, including:

- Secure guest account
- Application Layer firewall
- Application signing
- Sandboxes
- Full disk encryption
- Library randomization

Apple's intention was to provide for greater pre-emption of attacks as well as resiliency to them, and the company has continued its development of better security throughout OS X's lifecycle. Currently on version 10.8 (Mountain Lion), which was released in 2012, the platform includes many of the same features as previous iterations, with improved administrative controls and privacy options.

One of the biggest security challenges is conceptually not unique to Macintosh, but with the increased integration of Apple's iCloud services in OS X (not to mention iOS in the mobile space), there is a greater risk of files being moved from local to network

storage. Users expect access to data from any device, anywhere at any time, and that convenience makes all lost or stolen laptops, tablets, and phones significant risks.

While the MacOS's fortunes have waxed and waned and waxed again, there still are fewer security incidents reported because the greater popularity of other platforms make it easier to learn about them, because there are more potential targets, and because one single attack can affect more systems. Put another way, there are fewer attacks on Macs because the hacker community is not so familiar with them, and there are fewer attractive targets worth the opportunity cost. Still, it is just as important to keep the version of MacOS up to date as with other operating systems.

25.6 CONCLUSION. A good administrator can secure almost any NOS, although no NOS is secure initially. The network administrator needs continuous vigilance and monitoring, while recognizing that the operating system is only a part of the overall security plan for the LAN and network services. Most network administrators, due to the nature of their job and training, focus exclusively on the computers attached to the LAN and to the LAN's operating system and software. Unfortunately, this approach is too narrow in its scope. Personal firewall software also might be employed to protect individual systems against attack, but almost all of these products are oriented toward IP-based attacks and miss attacks that employ the NOS's native operating system.

Routers, network firewalls, and proxy servers are essential for protecting LAN systems from attack by an external source. The network administrator also must provide tools to protect servers and workstations from other users on the LAN.

25.7 FURTHER READING

This section lists some books, articles, and Websites that cover the issues addressed in this chapter. Administrators should monitor vendor, operating system, and security Websites that will have up-to-date, timely additional information.

Anonymous. *Maximum Security*, 4th ed. Indianapolis, IN: SAMS: 2003.

Bathurst, R., R. Rogers, and A. Ghassemlouei, *The Hacker's Guide to OS X*. Amsterdam: Elsevier, 2013.

Brenton, C., and C. Hunt. *Mastering Network Security*, 2nd ed. San Francisco: SYBEX, 2002.

Edwards, M. J. *Internet Security with Windows NT*. Loveland, CO: Duke Press, 1998.

Fraser, B., ed. *Site Security Handbook*. RFC 2196/FYI 8. September 1997. www.ietf.org/rfc/rfc2196.txt

Garfinkel, S., G. Spafford, and A. Schwartz. *Practical Unix and Internet Security*, 3rd ed. Sebastopol, CA: O'Reilly & Associates, 2003.

IEEE Working Group for WLAN Standards. http://grouper.ieee.org /groups/802/11/

The Information Warfare Site (IWS). Rainbow Series Library page. www.iwar.org.uk/comsec/resources/standards/rainbow/rainbow.html

Landau, T. *Sad Macs, Bombs, and Other Disasters*, 4th ed. Berkeley, CA: Peachpit Press, 2000.

L0pht Heavy Industries. Overview of AntiSniff. 1999. http://blockyourid.com/~gbpprorg/l0pht/antisniff/overview-2.html

Mann, S., and E. L. Mitchell. *Linux System Security: The Administrator's Guide to Open Source Security Tools*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 2002.

McClure, S., J. Scambray, and G. Kurtz. *Hacking Exposed: Network Security Secret & Solutions*, 6th ed. Berkeley, CA: Osborne/McGraw-Hill, 2009.

McNamara, J. The Complete, Unofficial TEMPEST Information Page. October 2, 2000. www.jammed.com/~jwa/tempest.html

Payne, W. H., and T. Sheldon. *The Complete Reference to Netware 5*. New York: McGraw-Hill, 1999.

Thurrott, P., and R. Rivera. *Windows 8 Secrets*. Indianapolis, IN: John Wiley & Sons, 2012.

University of Connecticut. (2004, June 17). IT Policies, Standards, and Guidelines: LAN Security Guidelines. University Information Technology Services Website. June 17, 2004. http://uits.uconn.edu/?page_id=499

Wi-Fi Alliance. WiFi Alliance home page. 2013 www.wi-fi.org

25.8 NOTES

1. See www.ietf.org/rfc/rfc2196.txt
2. See www.asystematics.com/asysus/c3-java-saf.htm, www.oracle.com/technetwork/java/index-jsp-141438.html, www.ehow.com/how_7160897_run-activex-controls-plugins.html, and http://windows.microsoft.com/en-US/windows-vista/Should-I-install-ActiveX-controls
3. For details of software-based TEMPEST, see Markus G. Kuhn and Ross J. Anderson, "Soft Tempest: Hidden Data Transmission Using Electromagnetic Emanations," 1998; www.cl.cam.ac.uk/~mgk25/ih98-tempest.pdf
4. See http://kl2217.wordpress.com/2009/07/31/ieee-802-11-standards-comparison/ and http://janmagnet.files.wordpress.com/2008/07/comparison-ieee-802-standards.pdf for WLAN physical layer details.
5. See www.wi-fi.org/files/wp_9_WPA-WPA2%20Implementation_2-27-05.pdf for WPA and WPA2 security details.
6. See www.securityinfowatch.com/directory for SecurityInfoWatch.com's Buyers Guide of Computer Security Products.
7. Microsoft came under fire for its original implementation of Kerberos in Windows 2000. Microsoft's version used a piece of proprietary data in the authentication process, making it impossible for third-party Kerberos servers to work with Windows. The company has since released interoperability information.
8. Ed Bott, "What Triggers User Account Control Prompts?" February 2, 2007; www.edbott.com/weblog/2007/02/what-triggers-user-account-control-prompts
9. Alec Klein and Ellen Nakashima, "For Windows Vista Security, Microsoft Called in Pros," *Washington Post*, January 9, 2007; www.washingtonpost.com/wp-dyn/content/article/2007/01/08/AR2007010801352.html

CHAPTER 26

GATEWAY SECURITY DEVICES

Justin Opatrny

26.1 INTRODUCTION 26·2

26.1.1 Business Requirements Outpacing Security 26·2
26.1.2 Demand-Driven Processing 26·3
26.1.3 Ubiquitous Mobility 26·3
26.1.4 Regulatory and Industry Compliance 26·3

26.2 BASIC CONCEPTS AND TERMINOLOGY 26·4

26.2.1 General Capabilities 26·4
26.2.2 Unified Threat Management 26·4
26.2.3 Next-Generation Firewall 26·4
26.2.4 Web Application Firewalls 26·4
26.2.5 Firewall Architectures Changing 26·4
26.2.6 Packet Filtering 26·5
26.2.7 Stateful Inspection 26·5
26.2.8 Application Layer Gateway 26·6
26.2.9 Current Gateway Security Devices 26·6
26.2.10 Host Environment Context 26·6
26.2.11 Firewall Platforms 26·7
26.2.12 Routing Devices 26·7
26.2.13 Appliances 26·7
26.2.14 Virtualization 26·8
26.2.15 Host-Based Agents 26·9

26.3 NETWORK-SECURITY MECHANISMS 26·9

26.3.1 Allowed Paths 26·9
26.3.2 Tunneling 26·10
26.3.3 Anti-Spoofing 26·10
26.3.4 Network Address Translation 26·11
26.3.5 Intrusion Detection 26·11
26.3.6 Intrusion Prevention/ Response 26·11
26.3.7 Encryption 26·13
26.3.8 Identity-Based Enforcement 26·14
26.3.9 Complex Protocol Engines 26·14
26.3.10 Content Control and Data Leakage Prevention 26·15
26.3.11 IPv6 26·16
26.3.12 Additional Considerations 26·17

26.4 DEPLOYMENT 26·18

26.4.1 Zoned Architecture 26·18
26.4.2 GSD Positioning 26·19

26.5 MANAGEMENT AND MONITORING STRATEGIES 26·20

26.5.1 Monitoring 26·20
26.5.2 Policy 26·21

26.6 MAINTENANCE 26·22

26.6.1 Patching 26·22
26.6.2 Pattern Updates 26·22
26.6.3 Logging and Alerting 26·23

26.6.4 Secure Baseline Configuration 26·23
26.6.5 Default Configurations 26·23
26.6.6 Harden the Device 26·24

26.7 DISASTER RECOVERY 26·24
26.7.1 Fail-Over/High Availability 26·24
26.7.2 Load-Balancing 26·24
26.7.3 Backup/Restore 26·24

26.8 NETWORK-SECURITY DEVICE EVALUATION 26·25
26.8.1 Current Infrastructure Limitations 26·25
26.8.2 New Infrastructure Requirements 26·25
26.8.3 Performance 26·26
26.8.4 Throughput 26·26
26.8.5 Implications of Added Features 26·26
26.8.6 Management 26·27
26.8.7 Logging and Alerting 26·28
26.8.8 Auditability 26·29
26.8.9 Disaster Recovery 26·29
26.8.10 Usability 26·30
26.8.11 Learning Curve 26·30
26.8.12 Features 26·30
26.8.13 Price 26·31
26.8.14 Vendor Considerations 26·32
26.8.15 Managed Security Service Providers 26·33

26.9 CONCLUDING REMARKS 26·34

26.10 FURTHER READING 26·35

26.1 INTRODUCTION. Once considered sufficient to protect an entire organization from external threats, the firewall is still perhaps the most recognized and deployed network-security devices for Internet-connected operations. However, earlier firewall generations made security decisions with little contextual support other than the origin and destination of the packets traversing a particular allowed path.

As communications capabilities and functionality demands increased, so too did the firewall's need to inspect and enforce allowed paths using more complex protocols and require ever-increasing throughput. This evolution transformed the firewall into a true gateway security device (GSD)—able to provide allowed path enforcement using a combination of techniques which once required additional security devices to accomplish.

Properly selected and deployed GSDs are one security layer designed to handle these increasingly complex scenarios. The GSD is effective only with a full understating of the capabilities and limitations—both operational and failure conditions—of consolidating multiple security functions into a single device. By providing allowed path enforcement more intelligently and accurately, combined with the added rigor of genuinely understanding expected network flows, GSDs provide sufficient additional defense-in-depth layers throughout the organization.

Although this chapter focuses on the GSD as a combined security device, the concepts covered are useful for understanding and evaluating the functionality of individual network-security devices. Every organization must make risk and performance decisions by weighing this approach against maintaining independent devices that focus on a particular security function.

26.1.1 Business Requirements Outpacing Security. Technological advancement continues to transform an enterprise's ability to manage the data life cycle. The pervasiveness of mobile devices and move toward cloud-computing resources that are no longer solely controlled or consumed by the organization continues to redefine the perimeter. Users are increasing demands for unfettered access to corporate data from anywhere on any device (including those not controlled by the organization).

This dynamic environment increases the need for layered security architectures with deeper awareness of content and context.

26.1.2 Demand-Driven Processing. Enterprises not possessing the requisite internal human or technological expertise to achieve the organization's goals for information technology (IT) have long looked to outsourced solutions to meet their needs. Software as a Service (SaaS) provides offerings such as productivity applications, collaboration, and email (e.g., Microsoft Office365) and customer relationship management (CRM) (e.g., salesforce.com). Infrastructure- and Platform-as-a-Service (IAAS and PAAS, respectively) provide on-demand storage and computing (e.g., Amazon Web Services—S3/EC2 or Rackspace Open Cloud). Outsourced offerings continue to mature and redefine how enterprises develop, manage, and present their information.

Virtualization technology continues to provide opportunities to use internal processing capabilities more efficiently. Although providing better performance, it can also diminish security in that existing offerings are not necessarily as mature as dedicated security infrastructure. When internal and Internet facing systems must operate on the same virtual architecture, this creates additional risk as systems now interact at the hypervisor and virtual switch (vSwitch) level where traditional network-security devices are unable to inspect and enforce traffic at this level.

See Chapter 68 in this *Handbook* for more details about security and outsourcing.

26.1.3 Ubiquitous Mobility. Today's business climate demands the ability for employees to work from anywhere, and this need for mobility and flexibility continues a significant shift in how organizations define and protect their perimeters. Employees may use a variety of systems whether at work, home, or on the road. The level of access and functionality required extends well beyond email into enterprise applications and data. As functional mobility stretches from company-owned to personal devices, organizations must have a method to ensure a compromised mobile device does not weaken the existing internal or outsourced security controls.

26.1.4 Regulatory and Industry Compliance. Federal regulatory requirements continue to have a significant impact on how organizations manage risk through data protection, retention, and privacy activities—many having significant auditing and reporting requirements.

- The Sarbanes-Oxley Act (SOX) focuses on controls and procedures designed to preserve the integrity of publically traded organizations' financial reporting.
- The Gramm-Leach-Bliley Act (GLBA)—specific to financial institutions—concentrates on the protection of customer data and privacy.
- The Health Information Portability and Accountability Act (HIPAA) requires protection of individually identifiable health information such as personally identifiable information (PII) and protected health information (PHI).

Industry-specific requirements continue to appear and evolve in an effort to address the minimum security requirements to operate with or within a specific industry. The Payment Card Industry Data Security Standard (PCI DSS) establishes baseline requirements for the protection of cardholder data during processing, transmission, and storage. PCI DSS requires organizations to determine their level of involvement

with the cardholder data. Once established, this determines which requirements are necessary to demonstrate compliance with the standard.

Each critical infrastructure sector—as defined in the United States by the Department of Homeland Security—has a Sector-Specific Plan (SSP), which provides representative organizations and agencies with the risk-management tools and strategies specific to the protection of each industry. One of the more mature programs is North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP). NERC CIP addresses multiple physical and digital security elements of the North American power system, including gateway protection through the establishment of an electronic security perimeter.

See Chapter 64 of this *Handbook* for more details on GLBA, SOX, and PCI DSS.

26.2 BASIC CONCEPTS AND TERMINOLOGY. Greater demands for mobility, new models for business interaction, and leveraging Internet-based processing capabilities continue to force the evolution of traditional perimeter protections.

26.2.1 General Capabilities. With the substantial processing requirements of specialized network security systems (e.g., IPS or anti-spam), organizations traditionally architected security infrastructures that leveraged dedicated solutions/devices for each function. However, current generation GSD processing capabilities provide the opportunity to combine many of these once dedicated systems into a single device.

26.2.2 Unified Threat Management. Unified threat management (UTM) combines elements such as anti-malware, anti-spam, IDS/IPS, VPN, application proxy, and content filtering—transforming the firewall into the original iteration of the GSD. These added capabilities allow the UTM to provide greater control and inspection at the application layer. However, typically there was only marginal management and performance integration between feature sets.

26.2.3 Next-Generation Firewall. The next-generation firewall (NGFW) is the latest evolution in stated capabilities that complements and surpasses those of the UTM. Instead of just bolting multiple security technologies on top of one another, the NGFW provides tighter integration of each level of security. These new capabilities include greater protocol awareness and more granular allowed path enforcement. The NGFW is able to profile protocols regardless of port chosen. This increases the ability to detect deceptive behavior such as encrypted payloads over protocols that would not normally use encryption. This generation of GSD also has the ability to adjust policy dynamically—extending protection to other parts of the security infrastructure.

26.2.4 Web Application Firewalls. GSDs provide additional capabilities to inspect and enforce allowed paths for Web-based communications. However, the complexity of the current and next-generation Web protocols may outpace the security provided by this device. The Web application firewall (WAF) provides more robust HTTP protocol inspection capabilities. The WAF provides customizable rules to protect against common payload-based attacks such as SQL injection, XSS (cross-site scripting), and command injection. This platform can also serve as a virtual patch for Web applications—legacy (no patches possible) or ones awaiting unreleased patches.

26.2.5 Firewall Architectures Changing. As security vendors work to keep pace with these changes with more functionality and higher performance, it is the

customer's responsibility to understand the advantages and disadvantages each proposed protection solution. Such analysis provides the necessary insight for the organization to deploy the most appropriate architecture to meet the necessary security and performance requirements.

26.2.6 Packet Filtering. Routing devices gave rise to first-generation firewalling capabilities. Packet filtering is a set of explicit rules describing the allowed paths network traffic may travel. The rules, in the form of an access control list (ACL), independently evaluate one or more portions of each packet's header to make the allowed path decision.

The packet filter acts on each packet as an individual entity without respect of the packet(s) that come before or after. Although this method provides the security with the least overhead of other firewall architectures, it is vulnerable to several network and transport layer attacks.

Internet Protocol (IP) *spoofing* is specially crafting a packet in an effort to deceive the router into accepting traffic that appears legitimate. The attacker will configure a packet to look like it originated from the internal network even though it is coming into the external interface. A *Land* attack sends a spoofed TCP SYN packet to a host with the source and destination IP and port being equal and can cause a denial of service on a vulnerable network stack. The *Teardrop* attack intentionally fragments a packet and manipulates the fragment offset where the preceding fragment's offset overlaps with the offset of the next fragment. When the receiving system reassembles the overlapping fragments, the network stack will crash, causing a denial of service.

Although both the Land and Teardrop attacks require a vulnerable end-device, initial packet filters did not have contextual understanding to stop such attacks. Mitigation for the Teardrop attack includes implementing packet reassembly—recombining the fragments into the original packet to ensure the recombined packet does not violate basic IP packet specifications—before forwarding to its next hop.

For details of these and other denial-of-service attacks, see Chapter 18 in this *Handbook*.

Current generations GSDs use stateful inspection (covered next) in an effort to overcome many of the limitations of packet filtering. However, many nonsecurity-related devices use packet filters to provide basic protection of their administrative interfaces and monitoring functions.

26.2.7 Stateful Inspection. Stateful inspection identifies and tracks additional parameters within each packet, adding context by representing flows along allowed paths as network connections instead of individual packets. Even though IP and User Datagram Protocol (UDP) are connectionless protocols, the firewall creates a virtual connection to emulate its connection-oriented counterparts. Although the specific parameters tracked in the state table vary across GSD vendors, the table resides in memory for more efficient processing.

Since a packet filter inspects packets individually, it does not affect firewalls in a load-balanced or fail-over architecture. However, to ensure sustained connections in high availability architectures, stateful inspection must establish and maintain a synchronized copy of the state table in each of its high availability members. Although some vendors still maintain that a direct serial connection is the most reliable connection method, highly available solutions are unlikely to be in a physical proximity necessary for serial connection and typically take advantage of existing network connections to maintain the state synchronization.

26.2.8 Application Layer Gateway. As threats became more sophisticated across existing allowed paths, firewalls added application-layer gateway (ALG) protection. Each application-specific proxy (e.g., HTTP, RPC, FTP, etc.) acts as the allowed path broker for the connection, creating a separate, backend connection to the other host. Although the full connection is actually two independent sessions, this duality is transparent to the requester and the server. Being a brokered connection adds an additional layer of defense by re-writing potentially malformed requests, validating/enforcing protocol compliance, and not allowing direct communication between client and server.

The detailed payload analysis increases the amount of time the firewall spends evaluating each packet potentially reducing throughput for other network traffic and increasing the need for sufficient processing resource. The ALG may also create issues such as poor performance for protocols requiring minimal overhead in addition to unexpected compatibility issues due to vendor-specific protocol implementations. Active review and management will help to minimize the ALG's impact on operations.

26.2.9 Current Gateway Security Devices. The firewall evolved yet again as attacks continued to increase in complexity and scale. Stateful filtering and application layer gateway functionality being effective at protecting against specific known attack vectors. However, the sophistication of the newer generation attacks quickly revealed those two features were unable to provide the necessary level of protection. The additional capabilities integrated into the firewall gave rise to the GSD.

The initial iteration of GSD is known as unified threat management (UTM). UTM is the consolidation of the firewall with multiple additional security platforms (e.g., network intrusion-prevention systems, content filter, anti-spam, etc.) into a single device. UTM vendors that only had a strong presence in one or two of the protection measures had to either integrate a third-party security product or build their own offering to fill out the full protection suite. This was also an opportunity to consolidate management and monitoring capabilities to create a unified platform. However, the UTM would prove somewhat inefficient, as each packet travels serially through the different security engines, severely impacting performance as the number of protections enabled increases.

The most current generation of GSDs—known as next generation firewalls (NGFWs)—provide greater application detection, awareness, and enforcement in addition to tighter integration of all security layers. Increasing the processing and inspection capabilities alone were not enough to meet the demands of the current generation application protocols. These protocols are more complex in form and function, even though they may follow the traditional allowed path such as HTTP. The NGFW must also be able to decode and understand the application's inner workings regardless of port and enforce on all or subsets of the protocol. NGFWs process packets though selected security engines in parallel to increase performance and focus the protection needs to the specific type of traffic detected.

26.2.10 Host Environment Context. Although this chapter focuses on gateway security devices from a network perspective, comparable host-based protections warrant coverage, as they are an important defense-in-depth component and have certain capabilities that the network-based GSDs cannot provide. With mobility and ubiquitous information access being central themes surrounding today's personal and business environments, the host-level protection is an ever-present fixture. The host's

movement from relying on the network security measures to local context and environmentally aware security measures will provide more robust and flexible security wherever the host goes.

Host-based security—aside from basic firewalling—has additional complexity and considerations due to the protection requirements of the additional elements (software, services, etc.) running. The host's ability to understand its local environment allows for more granular protection and visibility. Instead of focusing purely on port-based network access, the increase environmental context defines what applications and services can: access or receive network traffic, access or change other services, or execute code in a virtual machine (VM) to verify expect behavior(s).

26.2.11 Firewall Platforms. GSDs continue to advance in form, function, and contextual understanding to provide allowed path enforcement at all layers of the network stack. Malicious actors continue to develop increasingly targeted and effective code that travels on an otherwise allowed path. These more detailed network inspection requirements dictates that every organization select the most appropriate platform to meet its security and performance needs.

26.2.12 Routing Devices. As the need to transmit data within a network and beyond increased, routing device were adapted to provide native and then modular security services.

26.2.12.1 Access Control Lists. Routing device ACLs typically provide rudimentary allowed path enforcement with minimal performance impact. ACLs permit or deny individual network packets based on a combination of parameters from the header such as source and destination IP addresses and/or TCP/UDP port.

Routing device–based security services matured to provide enhanced capabilities based on their contextual understanding of the network. Routing devices can support several stateful inspection techniques, the most rudimentary being the ability to track connections based solely on the acknowledgement (ACK) flags of an established connection. However, this method is more susceptible to spoofed packets directed at the intended target than traditional stateful inspection based on a table of established connections. Routing platforms can also provide anti-spoofing capabilities based on network topology. When enabled, the routing device will make forwarding decisions based on its understanding of connected and learned routes and only allow source IPs to originate from the appropriate interface.

26.2.12.2 Hardware Modules. With bandwidth increasing into the tens to hundreds of Gigabits per second (and beyond), routing device manufacturer needed to offload the security functions to a dedicated vendor-specific module (or line card). The module provides logical protection for one or more allowed paths while using the backplane to achieve higher data rates and provide protection based on one or more of the firewall architectures described above. Some routing vendors also support modules—similar to a blade server—allowing other security vendors to integrate their technology deeper into the routing infrastructure.

26.2.13 Appliances. The GSD appliance emerged as a way to shed complex security decisions from routing devices onto a dedicated security platform. Appliance form factors, capabilities, and limitations vary among GSD vendors with the GSD application integrating into the chosen hardware platform.

26.2.13.1 Proprietary. A proprietary hardware-based GSD appliance manufacturer is able to establish and retain tight controls over all aspects of the platform. The hardware is purpose-built to interoperate with the GSD software and provides different levels of performance and functionality. High-performance components such as application-specific instruction sets are typically proprietary. However, the vendor will balance the cost effectiveness of the platform by choosing commercial-off-the-shelf (COTS) hardware for the utility components (e.g., network interfaces and RAM). The accompanying proprietary, hardware-specific operating system (OS) allows the vendor to maintain tight controls on interoperability, configuration, and software revisions.

26.2.13.2 Hardware Independent. Hardware independent GSD appliances transfer responsibility to the customer to procure hardware that meets a specific set of requirements while allowing the GSD vendor to retain responsibility for developing the GSD application and/or OS. The GSD uses either a customized version of an OS (typically Linux or Berkeley Software Distribution [BSD]) or installs directly onto a customer-provided base OS.

A GSD application installed on a proprietary (or premodified) OS reduces native security exposures through vendor custom hardening of OS mechanisms to remove nonessential features and functionality while maintaining the tight controls described in the section above.

The GSD application may also be able to install on specific versions of the customer's native OS choice, providing the customer the most flexibility in choosing a GSD appliance. The GSD application will make some hardening changes but leaves much of the operating system management to the customer. This increases the risk of insecure and/or underperforming implementations due to the GSD vendor not fully controlling OS configurations, patches, and interoperability.

26.2.13.3 Soft Appliance. The soft appliance augments the hardware independent/proprietary OS GSD into a bootable platform that can run on consumer-grade equipment. This platform is more suited for point-in-time protection needs, such as main GSD device failure, fast-turn deployments, and testing. Before deploying a soft appliance, the customer must temper the potential "flexibility" of this solution with other considerations such as manageability (configuration, logging, etc.), scalability, and required functionality.

26.2.13.4 Embedded Appliance. In addition to high-end, scalable devices, GSD vendors also recognized the opportunity to service other markets such as providing specialized capabilities on platforms scaled to meet the needs of consumers and small-to-medium businesses (SMBs).

SMBs are less likely to have sufficient funding or dedicated security staff to purchase, manage, and monitor multiple security devices. The SMB-grade GSDs provide a more economical all-in-one device, due to its smaller form factor, simplified user interface, and lower performance requirements. In addition to providing the standard wired and/or wireless connectivity options, the device typically provides a full complement of UTM security functionality.

26.2.14 Virtualization. GSD vendors now provide solutions ranging from leveraging the virtual environment to provide independent network security capabilities to embedding granular protections at the hypervisor level. Instead of requiring

a full installation, some GSD vendors will provide a preinstalled virtual appliance that minimizes total deployment time. It is important to ensure that the virtualized GSD maintains its priority in obtaining and sustaining resources from the VM instance.

A VM-capable GSD acts as an independent network security entity. Instead of having dedicated hardware, the GSD leverages the hypervisor's hardware abstraction layer to gain access to the necessary resources. This implementation provides hardware appliance-like capabilities to control the allowed paths based on what networks route through the GSD VM instance.

A VM-embedded GSD virtual appliance integrates with the hypervisor to gain low-level access to the entire VM host/cluster. The GSD now has visibility into previously inaccessible communication paths and is able to conduct allowed path enforcement for both inter- and intra-VM communications.

26.2.15 Host-Based Agents. As mobile devices demand greater internal network access—including data to be stored local to the device—so did the requirement for those host-based protections to be able to minimize the additional risk of allowing such access. Antivirus capabilities, although arguably only marginally effective against new and emerging threats, still maintain a necessary presence to filter known attacks vectors. Firewalling capabilities at the network and application level reduce the host's network-facing threat profile. The Host Intrusion Prevention System (HIPS) provides protection against known and potentially unknown threats because of its ability to inspect up to the application layer.

When network-security devices fail to detect a threat, host-based protection has additional measures to detect local anomalous behavior due to its increased contextual awareness of how the host—the applications and the operating system—should behave, thus allowing the host an advantage in detecting emerging threats. However, organizations should not underestimate the administrative overhead necessary to manage and monitor this protection.

See Chapter 27 of this *Handbook* for more details on host-based protections.

26.3 NETWORK-SECURITY MECHANISMS. Network security, once synonymous with the firewall, continues its growth in both scope and depth. Government and organizations are requiring security to be baked into the systems and solutions they produce or purchase.

It is entirely possible for an attack to originate from a system not coinciding with a direct perimeter attack. These attack vectors include likely exposures, such as an employee mobile device infected with malicious code while traveling; the new USB flash drive an employee receives at a conference; or a vendor providing legitimate support that plugs his/her laptop into the internal network.

Those who are selecting network security systems must be able to decipher the hyperbole from reality to ensure the device(s) deployed realistically address(es) the specific risks pertinent to the organization. Gateway security devices maintain their basic allowed path control heritage but now are far more advanced in their ability to integrate protection within the payload and beyond.

26.3.1 Allowed Paths. Network-security devices provide allowed-path protections to ensure network traffic only flows in expected and intended ways. Unfortunately, this base level of enforcement is insufficient because application-layer attacks (e.g., client-side browser exploits) leverage these traditional allowed paths. Regulatory and compliance security requirements are also driving the deployment of more

capable allowed-path enforcement mechanisms deeper into the internal network to protect intellectual property, customer/partner data, and industrial control systems (ICSs).

To manage allowed-path risk effectively, organizations must define and baseline their expected traffic flows both internally and externally. The baseline is vitally important when the only indication of compromise may be an unusual—though apparently benign—interaction between systems that do not typically interact with one another.

26.3.2 Tunneling. Tunneled access into the internal network for the purposes of remote access, information sharing, and commerce is essential for most organizations. These tunnels, typically in the form of a virtual private network (VPN), use encryption to protect the confidentiality of the transferred data. If the GSD does not terminate the tunnel, it typically only provides network and transport layer allowed path enforcement through to the termination endpoint. This scenario fragments the inspection of the encapsulated traffic, requiring the termination endpoint itself or another security device post-decryption to provide additional protection through allowed path enforcement and application layer inspection.

See Chapter 32 for more information on GSD VPN capabilities

26.3.3 Anti-Spoofing. One of the most readily exploited vulnerabilities of the IP protocol is *spoofing*. Due to this inherent weakness in the IPv4 protocol, a host's network stack is capable of producing packets that can appear to originate from any IP address and/or port chosen. This attack is an effort to deceive the receiving host and/or any network protection in between into believing the originating host is following an allowed path. Due to the simplicity of executing spoofing attacks, GSD vendors developed and implemented several effective methods to thwart these types of attacks. The network topology is a key element in the GSD's ability to determine what traffic should be originated from each of its physical and/or logical interfaces.

One common anti-spoofing method is explicit definition of the network topology on a per-interface basis. This provides an efficient method for the GSD to determine whether the packet's source IP address entering an interface matches what the GSD expects to be the source network. If not, the packet fails to meet the criteria and is dropped. Although simple and effective, the disadvantage is the manual management of the per-interface network definitions.

Another method is to learn the network topology by leveraging the routing table to make the same type of per-interface anti-spoofing decisions. The firewall goes through the same process of evaluating if the packet's source IP matches what should be originating from that interface. However, this does not require ongoing manual intervention when network changes occur, as the routing engine updates this information. Although this method creates less administrative overhead, two disadvantages include problems if the network uses asymmetric routing and/or falls prey to a route poisoning attack.

Both of these anti-spoofing measures are highly effective when evaluating whether network traffic is originating from the appropriate interface. However, the focus remains solely on preventing internal network address compromise. Since not all of the IP networks in existence are actually allocated and/or in use, these "bogon" networks are another method used to launch spoofing attacks. Since "bogon" network traffic is originating on the external interface, the GSD would not block this based on common anti-spoofing rules. Updated lists are readily available allowing the security administrator to enforce "bogon" protections. Some GSDs have a feature that automatically updates the list on a recurring basis, providing automated protection.

26.3.4 Network Address Translation. Network address translation (NAT) is a mechanism that maps internal (typically private RFC 1918) network addresses to one or more publicly routable IP addresses. This is an essential function in IPv4 due to the increasing shortage of publically routable address. Due to its ability to mask the internal network addressing, NAT quickly became thought of as a security mechanism. Realistically, NAT only ensures network traffic maintains their translations—but it has no inherent security capabilities such as payload inspection.

To ensure sustained connections in high availability architectures, NAT must establish and maintain a synchronized copy of the translation table with each high availability member. This mechanism follows the same serial or network connection used to maintain the state table.

26.3.5 Intrusion Detection. GSDs inherently provide detective controls at multiple layers. These distinct controls (e.g., firewall, intrusion detection/prevention system, etc.) provide greater visibility into the types and scopes of potential incidents when properly configured and monitored. The logging and alerting capabilities associated with each control are important and necessary mechanisms for security personnel to use to understand the current state of the perimeter. With this information, it is possible to develop actionable intelligence to detect and act on anomalous and/or malicious behavior before, during, and after an incident.

Responding to alerts and reviewing all of the information generated by the GSD can be extraordinarily time consuming. However, it is important that the security administrator take the necessary steps to ensure timely and accurate log reviews. Without such reviews, it is easy to overlook issues such as pre- and post-attack-related traffic and system misconfigurations. Logs also provide an excellent opportunity for GSD tuning, including rule-base optimization and reducing false positives (postverification).

Properly planned and configured alerting and logging mechanisms are essential. When done poorly, these mechanisms can overload the support staff with unnecessary or extraneous information, making them less efficient and effective at addressing actual security issues.

26.3.6 Intrusion Prevention/Response. Although logs and alerts are useful for detecting and investigating security incidents, these mechanisms are passive and cannot provide threat protection. Firewalls evolved to include automated active mechanisms to respond to security events and provide active protection.

See Chapter 27 of this *Handbook* for more details on intrusion detection and prevention.

26.3.6.1 Connection Termination. The essence of perimeter defense is to protect the internal network by blocking connection attempts that do not follow the established allowed path policy. One method for TCP-based protocols is to respond to the initiator (and sometimes the intended recipient) with a reset (RST) packet, severing the connection. The most common method—which works for connection-oriented and connectionless protocols—is to drop the packet without a response. As firewalls evolved into richer functioning GSDs, this protection grew to include blocking based on malicious activity occurring at the application layer.

Inversely, attackers can use connection termination as a method to determine where and what type of network-security devices are in use. This context may provide them additional insight on potential methods to bypass and/or adversely impact the security

platform. Though this situation is unavoidable no matter what termination mechanism is employed, it is another reason to apply defenses in layers.

26.3.6.2 Adaptive Threat Mitigation. Although the GSD's policy provides the main enforcement mechanism, vendors quickly realized that this was insufficient, as attackers continually develop additional methods to thwart these protections. GSDs include several enhanced capabilities that allow the device to adapt to specific threat conditions and provide targeted prevention.

Since attackers employ methods to change attack sources to avoid detection, and it is not always practical to follow a rigid whitelisting (least privilege) strategy, organizations typically opt to use blacklists as an additional level of protection. Some GSDs use dynamic object groups to update these lists on a predetermined schedule helping minimize administrative overhead.

The use of thresholds provides the GSD a more efficient method to address consistent attacks that would otherwise follow the same policy-based evaluation for every attempt. One common threshold type is event quantity/type over a given set of time (e.g., total number of packets per second or x number of packets per second across multiple destination IP addresses). Once breached, the GSD will typically automatically block all traffic from the origin IP address for a predetermined time period. This type of adaptive technique requires additional processing resources and requires tuning to ensure this does not impact legitimate network flows.

Being primarily perimeter devices, GSDs typically have minimal visibility into the internal, trusted network. However, modern security infrastructures can leverage multiple threat detection methods across the network to provide more comprehensive protection. Once detected, the device can leverage the security management infrastructure to update other security device policies—automatically providing additional protection or containment. This capability is also integrating with host-based protections to drive protection deeper into the network. If a host detects a threat while on an untrusted network, the infrastructure would learn about this and adapt network protections before the device connects to an internal network.

26.3.6.3 System-Level Actions. As an added layer of defense, gateway security devices typically have built-in mechanisms to detect and respond to threats against the platform itself. Following an implicit deny model, the excepted behavior is to fail secure (e.g., shutdown the firewall) and not pass traffic. However, as this automatic response affects the functionality of the security device, it is important for the administrator to understand and test this functionality prior to deployment to ensure this response behavior is consistent with the protection and availability needs of the organization.

26.3.6.4 Application Inspection. The ability to protect allowed paths based on IP and/or port combinations provides marginal benefit when matched against the growing protocol complexity and configurability. As GSDs evolved, so did the ability to evaluate and enforce allowed paths based on the application layer. This capability ranges from basic protocol evaluations such as an RFC conformance check to full application layer gateways. Depending on the type of application layer protocol, a dedicated device may be necessary in high capacity, security, and/or reliability environments.

WAFs. Web application firewalls (WAFs) provide targeted enforcement within an HTTP(S) allowed path. WAFs can more readily detect and prevent common Web server attacks such as XSS and SQL injection and typically have added capabilities to protect

more complex Web 2.0 (e.g., AJAX) and Web services (e.g., SOAP or JSON) protocols. It is also possible to protect otherwise vulnerable Web servers from exploits by writing protection rules that detect and nullify the application layer attack traffic. The WAFs granular rule development is also better suited to protect custom Web applications.

Proxy Servers. The forwarding proxy enforces outgoing Internet connections by intercepting and controlling access based on criteria such as user ID, time of day, and/or whitelist/blacklist. Depending on the type and configuration, the forwarding proxy can either broker the outbound connection and make requests on behalf of the client or just enforce the particular allowed path. Classic forwarding proxies requires client-side configuration to enforced Internet-bound communications. A transparent proxy is a forwarding proxy that operates as a "bump-on-the-wire"—sitting in-line or off a span port—and typically does not require client-side configuration.

Conversely, a reverse proxy provides protection to inbound allowed paths. The inbound connection terminates at the proxy, and the proxy itself establishes the second half of the connection directly with the application server. Being a brokered connection, this adds an additional layer of defense for inbound traffic by rewriting potentially malformed requests and/or not allowing direct communication to the application server.

Application Identity. In addition to source and destination IP address, network firewalls typically rely on the destination port to make a final allowed path determination for common protocols (e.g., HTTP—TCP port 80). However, it is trivial to evade upper layer inspection (e.g., proxy or WAF) by changing the destination port to something that bypasses the application layer inspection but still follows an allowed path. Current generation security devices must be able to profile all network traffic based on application type instead of destination port.

Application identification must be able to dissect the traffic and understand the subsystems/protocols embedded within the main communication flow. For example, Skype not only inherently provides video conferencing but also file transfer and chat services. With add-ons, functionality can extend to remote desktop sharing/control. By having this detailed level of application understanding, it is possible for granular enforcement of one or more protocol subsets. This profiling also includes determining common protocols on unexpected ports and encrypted traffic over a common clear-text protocol—whether encrypted or not.

26.3.7 Encryption. Confidentiality protection for sensitive data-in-motion continues to grow due to the need to minimize the risk of compromising personal/personnel data and/or intellectual property. This creates a distinct need to provide allowed path enforcement when using encryption services to or through the GSD. Due to the computationally complexity, encryption protocols can significantly reduce GSD throughput without offloading encryption services to an add-on, on-board device, or dedicated external devices to meet the encryption needs.

See Chapter 7 of this *Handbook* for more details on encryption.

Inspection. The integration of encryption into network and application flows in essential to protect the confidentiality of the data-in-motion. Although not necessarily malicious, these encrypted sessions reduce the effectiveness of any network-security device(s) attempting to inspect the protected payloads. In addition to the typical VPN and HTTPS Web traffic, the GSD must also be able to identify encrypted communications regardless of port or service. The two main methods for inspecting encrypted traffic are direct termination and on-the-fly decryption.

In the first method, the GSD terminates the encrypted connection. This provides an opportunity to inspect the once tunneled communication. Once inspected and if

permitted, the GSD may pass the remained of the communication decrypted or re-encrypt the session to maintain the confidentiality protection to the endpoint.

In the second method, the GSD will use escrowed encryption keys to decrypt network flows passively. Besides meeting the additional performance requirements to conduct passive decryption, another possible weakness is the inability to provide fully synchronous responses to detected threat—potentially allowing some malicious traffic across an allowed path before the enforcement occurs.

VPN. A virtual private network (VPN) is an encapsulated network overlaying an existing set of physical and logical networks. A common VPN implementation is for a remote client connection—allowing an authorized user to access the internal network through an encrypted tunnel. Once established, VPN clients can access internal networks and/or systems while protecting the confidentiality of the connection. Site-to-site VPNs allow authorized remote locations to gain direct internal network access for multiple users over a single connection.

Since there is little to no control when mobile systems leave the confines of the internal network, the fidelity of the mobile device is a concern. The risk of compromise at every level—from physical to operating system to applications and beyond—can expose the internal network to increased risk during a VPN session. Without terminating VPN sessions directly on the GSD, it is still possible to increase visibility and additional allowed path enforcement once the unencrypted traffic leaves the VPN termination device—see the Deployment section for additional details and considerations.

See Chapter 32 of this *Handbook* for more details on VPNs.

Acceleration. Using Secure Sockets Layer (SSL) and Transport Layer Security (TLS) to protect Web-based traffic or IPsec to protect VPN traffic creates a potential barrier to comprehensive allowed path enforcement. Since the encrypted traffic potentially fragments allowed path enforcements, GSDs evolved to be capable of terminating these connections to add an additional layer of inspection. However, encryption/decryption is mathematically intensive and requires sufficient processing resources to ensure these processes do not impact GSD throughput. When it is necessary for the GSD to terminate or inspect encrypted traffic, the use of a hardware encryption module is a common method to off-load this processing burden and minimize impact to other traffic flowing through the GSD.

26.3.8 Identity-Based Enforcement. By integrating with enterprise directory services using mechanisms such as Lightweight Directory Access Protocol (LDAP) or Remote Access Dial-In User Service (RADIUS), GSDs are able to make identity-based decisions. Instead of allowing or denying access to an allowed path by source IP, the GSD can authorize access per user or based on group membership and include this authorization data in logs. This additional detail enhances reporting and auditing capabilities, providing more specific information about a particular connection well beyond just the IP address or DNS name of the host requesting allowed path traversal.

26.3.9 Complex Protocol Engines. Complex protocols, including those used to deliver feature rich Web applications, IP telephony, or specialized protocols, such as those used by industrial control systems (ICSs), are an intrinsic necessity for many organizations. Due to the ever-increasing need to share information and access regardless of mobile device form factor or location, these complex protocols are traversing the GSD. Even with successful detection of an application in-use, it is a continual

development race for the GSD to be able to understand and enforce the intricacies of each protocol.

Today's Web-based applications continually develop new functionality to meet customer demand. Social media Websites offer viable business models and may necessitate some or all users within an organization to access this service to conduct business. Due to ever-present personal elements (e.g., chat/comments, games, etc.), the GSDs must be able to dissect the connection and selectively permit some or all of the functionality offered by the host application.

IP telephony protocols such as session initiation protocol (SIP) and real-time protocol (RTP) are typically intolerant to latency and jitter—which are introducible by GSDs. ICS protocols introduce vendor-specific protocol dependencies/intricacies and are even less tolerant to network variations. The loss or slowness of ICS communications can be as severe as damage to equipment or death when impacting safety systems.

Although an internal or external network perimeter is a logical point of inspection, the GSD is not necessarily the appropriate device to conduct such detailed protocol inspections. Instead, these situations may warrant additional security layers using specialized devices for the application layer or specialty protocol enforcement.

26.3.10 Content Control and Data Leakage Prevention. Organizational policies establish the behavioral expectations of its employees. The Acceptable Use Policy (or similarly named) typically covers matters specific to use of technology resources; however, the policy itself cannot provide active protection. Content filtering (CF) bridges the gap between technical and nontechnical policy enforcement. This enforcement typically focuses on internal users attempting to access information resources outside of the organization. By filtering heavily used protocols such as HTTP and SMTP, the organization can minimize user access to information not fitting the organization's definition of business use.

This technology has several methods for inspecting, classifying, and filtering content. Inspection takes multiple forms, including residing in-line, out-of-band to monitor passing traffic without impacting the physical connection, or in conjunction with proxy redirection. Classification occurs through multiple methods, including vendor-defined categories, blacklists, and/or reputation scoring based on IP address and/or DNS name. This also allows the organization to choose to use a subset of the provide policy and/or selectively develop a whitelist to provide more granular enforcement. The filtering feature logs and/or severs the connections allowing the organization to customize its response to detected policy violations, including redirection to a "not for business use" Web page and/or automatic administrative and/or managerial notification of repeated violations.

See Chapter 31 of this *Handbook* for more details on content filtering.

26.3.10.1 Information Classification and Enforcement. An overlooked but important aspect of information assurance is the proper classification of information/data into categories. This allows the organization to determine the risk reduction efforts necessary to protect each information level/type. Similar to the classification process used by governments, identifying the sensitive information increases the ability to enforce not only who has access to the information but also where and how this information can be transferred.

Digital rights management (DRM) and data loss prevention (DLP) are two current generation technologies for addressing these information-specific risks. DRM

focuses on the definition and enforcement of data-at-rest—typically at the information repository level. GSDs may include DLP protections such as enforcing data-in-motion protections, including denying the transfer of files based on content/designation and/or requiring encryption for sensitive data.

26.3.10.2 Anti-malware. A contemporary attack vector is to target a commonly used Website/service and inject malware. Although the content filter may make an initial decision to allow the connection based on the specific site name, it may also have the added capability to intercept and block the malicious content before it ever reaches the host. This functionality makes it possible to evaluate different types of traffic such as SMTP attachments and HTTP downloads. Although this type of application layer protection can dramatically reduce the GSD throughput, it is possible to redirect the scanning to an additional module within the GSD with dedicated processing power or to a dedicated network anti-malware appliance.

26.3.10.3 Active Content. Active code used on HTTP connections—such as Silverlight, Java, AJAX, and Flash—provides a foundation for a rich Web experience but also increases security risk due to their application-focused functionality. Malicious email attachments—such as documents with embedded code—are common components of phishing and spam campaigns.

In conjunction with host-based protections, the GSD can decrease the risk of active code by preventing specific types from ever reaching the endpoint and/or scanning the active code in a sandbox prior to allowing the active code to traverse the allowed path.

See Chapter 17 of this *Handbook* for more details on mobile code.

26.3.10.4 Caching. Although the aggregate cost of bandwidth continues to fall, organizations continue to look for opportunities to manage these frequently congested circuits more efficiently. Proxy servers typically integrate caching mechanisms to store frequently used Internet content locally, increasing local throughput by reducing Internet-bound traffic. The GSD may also be able to provide similar caching services, given sufficient storage capacity and processing resources.

26.3.11 IPv6. The allocation of the last remaining available IPv4 address blocks to Regional Internet Registries (RIR) occurred in 2011—indicating the world is "out of IPv4 addresses." Although one may argue that statement is overly dramatic, mass migration to IPv6 is still occurring at a slower pace than one may expect, even though IPv6 was developed in the late 1990s. While organizations continue to develop their individual transition plans, it is likely that IPv6 is already in their environment. Many modern operating systems and devices have dual IPv4 and IPv6 network stacks, with some enabled by default. IPv6 and its associated transition technologies have specific security implications that the GSD must understand and enforce.

26.3.11.1 Perimeter Security Concerns.

Addressing. IPv6 is more flexible in its approach to dynamic addressing. Instead of solely relying on DHCP, an IPv6 device can address itself through stateless address autoconfiguration (SLAAC). The host uses a unique identifier (typically its own Message Authentication Code (MAC) address) in addition to the Neighbor Discovery (ND) protocol to complete the automatic addressing. Since there is no authentication

requirement, the GSD must prevent external devices from attempting to act as an internal router during the addressing process.

The significant increase of available addresses in any particular IPv6 network makes it infeasible to discover devices and network topology using traditional port scanning methodologies. By using the multicast listener discovery (MLD) protocol, an attacker can send a probe to the link-local multicast address (ff02::1) and listen for responses. The GSD must block this capability at the perimeter to prevent external devices from attempting to discover internal hosts and topologies.

Tunneling. Without ubiquitous end-to-end IPv6 connectivity, there are several IPv6 transition technologies (such as 6to4 and Teredo) that allow IPv6 capable systems to tunnel communication over legacy IPv4 networks. As with other tunneled traffic, to be effective, the GSD must not only be able to enforce the appropriate allowed path for the tunneled traffic but also inspect the encapsulated IPv6 packet.

IPv6 also has native support for IPsec (both AH and ESP). Configuration and use of IPsec in an IPv6 environment requires the same discipline in choosing and configuring cryptographic options as IPv4. It is also possible to use IPv4 IPsec to protect IPv6 transition technology tunnels as unencrypted tunnel sessions would otherwise be vulnerable to interception and/or manipulation.

Global Connectivity. Network address translation (NAT) is an essential element for IPv4 Internet connectivity due to the relatively "small" number of publically routable addresses. The size of the IPv6 address space (2^{128} addresses) provides for global IP interconnectivity without the need (or definition) of a NAT replacement. IPv6 devices will now communicate natively across the Internet—changing perimeter dynamics yet again. With IPv6 enabled, it is essential for GSD to enforce strict ingress *and* egress filtering.

Mobility. Mobile IPv6 (MIPv6) allows hosts to maintain access to the home network while physically roaming to other locations and uncontrolled networks. A user would be able to leave the office network and travel from one appointment to another without losing connectivity to the internal network. Since the device can move network to network without dropping the internal connection, MIPv6 also creates potential issues with stateful inspection if the GSD does not understand the protocol. If the organization chooses not to support MIPv6, the GSD should filter Type 2 Routing Header (RH2) packets.

26.3.12 Additional Considerations

26.3.12.1 Host Protection. Although firewalls and GSDs play a crucial role in protecting the overall network infrastructure, it is neither cost feasible nor possible to deploy these devices at all points inside the network. Individual hosts must have a way to protect themselves from threats independent of network-security devices. Unlike the network, a host has a contextual understanding of what the system can, is, and/or should be doing. Beyond the simple allow and deny functionality, contextual security measures can detect unexpected system configuration changes such as service changes or an application behaving in an unexpected manner.

When a host leaves the internal network, it becomes an extension of the network protection profile. By providing adequate local protections at the network and application levels, the mobile endpoint helps to mitigate potential issues upon reconnecting to the internal network. It is also crucial to determine the level of protection necessary and how each of these additional levels of security will affect system performance, management, and end-user impact.

26.3.12.2 Network. Hosts need to be able to determine the types and appropriateness of inbound and outbound traffic. These network restrictions may vary from the internal network to uncontrolled networks. In certain circumstances, all network traffic may be suspect and scrutinized further. For example, when on the internal network, the host allows most inbound and outbound traffic. If the host were on an uncontrolled network, it would not allow any nonestablished network flows. Simple host-based network protections are not enough; additional host protection mechanisms such as intrusion prevention can detect and stop network and application-based attacks.

26.3.12.3 Applications Access. The host's contextual awareness also helps dictate the ability for applications to execute as well as send and/or receive data. The goal is to ensure only the appropriate applications and/or services have network access. The host protection policy may allow applications to establish outbound connections, but never listen (nor accept nonestablished inbound packets). For example, an HTTP server uses a daemon to listen for connection attempts. The HTTP client will attempt to make a connection to the HTTP server daemon. By using a host protection mechanism, it would be possible to prevent one or both of these actions.

26.3.12.4 Hybrid Protections. The host intrusion prevention system (HIPS) functions similarly to a network intrusion prevention system (NIPS) by detecting known attack patterns and/or anomalous behaviors. Hybrid host protections build on the host's contextual awareness and provide the ability to monitor other unusual application level activity such as changes to binaries, service manipulation, and spawned listeners.

26.4 DEPLOYMENT

26.4.1 Zoned Architecture. In a contemporary interpretation of the screened subnet architecture, zones define different types and/or sensitivities of networks, applications, and/or services. This provides the ability to manage allowed paths at the macro level (per zone) in addition to the traditional allowed paths (specific hosts or services). As requirements drive security deeper into the network, the zoning concept is equally effective when used to manage and protect internal and external networks.

26.4.1.1 Perimeter Zones. The border router maintains the architecture's first line of defense against external attacks. The ACL(s) on this router should mirror the basic allowed-path configuration of the external (untrusted) firewall interface and provides several important benefits.

The GSD is able to operate at optimal efficiency, since traffic rejected based on border router's packet-filtering rules normally would never reach the firewall. This permits the firewall to focus, in terms of load, on protocol inspection. If, for example, the firewall receives a packet that should never have made it passed the border router's ACL, the firewall can assume that the router is not behaving normally. The firewall is then free to respond appropriately, with such actions as terminating all connections from a specific host.

26.4.1.2 External Service Zones. The necessity for mobility and accessibility places significant demands on Internet-facing systems in addition to increasing the administrative overhead of managing external access. The sections below provide only a sampling of possible external service zone architectures.

Utility. Instead of lumping systems such as Web, DNS, and email onto a single network, there may be an advantage by implementing zones. Utility servers such as DNS and email could logically be on the same network. Web servers typically demand greater bandwidth, and by using this concept, can protect the entire zone with a simple inbound access rule.

Extranet. Extranet systems create additional complexity since they provide the user interface, while internal systems may provide the relevant content. Zoning provides more flexibility by allowing external connections to reach the Extranet servers while providing those same servers access to internal resources.

VPN. VPN networks are also an opportunity to use zoning. Since the VPN connection device must be Internet facing, this requires two different networks connected to the firewall. The external, Internet-facing (VPN untrusted) interface would only allow a few protocols for inbound and outbound encrypted traffic. The second network (VPN trusted) is for unencrypted network traffic moving to and from the internal network. This architecture also provides extra internal network protection in the event of a compromise of the VPN device as well as creating a traffic inspection point that is unimpeded by encryption.

26.4.1.3 Internal Service Zones.

With the increasing prevalence of non-company-owned assets and more complex data flows, organizations continue to look for better methods to protect internal networks. The sections below provide only a sampling of possible internal service zone architectures.

Administrative and Monitoring Systems. Organizations typically limit direct access to network and security devices. By requiring all administrative and monitoring traffic to originate from a specific host and/or network, a zone can effectively reduce each device's threat profile. This zone would not only allow minimal noninitiated inbound traffic but would also limit outbound connections to the managed and monitored systems.

High-Value Systems. Organizations relying heavily on intellectual property and/or other protect data types (e.g., personal and financial information) have an intrinsic need to provide higher level protections to protect their investments. Zoning provides an additional layer of protection by minimizing unnecessary information flows to/from these high-value systems.

Industrial Control Systems. Incidents such as Stuxnet and Night Dragon are stark reminders that even though certain system types are considered extremely complex and less accessible, it is only a matter of time before successful compromise is possible. Manufacturing organizations with significant investments in industrial control and supervisory control and data acquisition (SCADA) systems can (or have requirements to) isolate these environments as an additional layer of protection and/or compliance.

26.4.2 GSD Positioning.

The increased use of encrypted protocols such as SSL/TLS and IPsec can blind network protections. Certain GSDs have the ability to terminate encrypted sessions, though the increased processing and bandwidth requirements may exceed the limits of the device. If concerned, the security architecture should deploy appropriate countermeasures at strategic locations that avoid encrypted traffic. This way, the GSD can focus on its primary role of detecting and preventing malicious activity.

26.4.2.1 Inline.

Placing the GSD inline creates a choke point for active enforcement on all network traffic that flows through it. When a malicious packet enters

the GSD, protocol analysis will detect the anomaly and will not allow it to flow out the other interface. Although bandwidth limitations are a typical concern, improperly configured inline devices may also present a denial-of-service condition. With proper infrastructure planning and deployment, it is possible to minimize these risks.

26.4.2.2 Controlling Encrypted Traffic. Since mobile devices frequently venture outside of the controlled network, one logical place to evaluate traffic is on the unencrypted side of the connection. This may be on the backside of a SSL terminator (in some cases on the server itself) or on the unencrypted side of a VPN connection. The second option is to use the GSD to inspect (e.g., termination and/or passive decryption) the traffic before forwarding the traffic onto its final destination. This level of functionality requires substantial processing resources but may be a necessity if the security layers downstream are insufficient.

26.5 MANAGEMENT AND MONITORING STRATEGIES. Regardless of vendor claims, network-security devices are never a plug-and-play endeavor. It is essential to take additional steps to define the security requirements for managing and monitoring GSD components. This approach helps ensure a well-rounded security posture.

26.5.1 Monitoring. Firewalls and GSDs provide complex functionality; monitoring such systems must go beyond just verifying system availability and cover device health, availability, and integrity.

26.5.1.1 Health. Metrics such as processor utilization, available RAM, and number of connections all have an impact on overall functionality. A centralized management console may provide the ability to monitor and alert on these metrics. If this functionality is unavailable, it may be necessary to use monitoring protocols such as Simple Network Management Protocol (SNMP) and/or Remote Monitoring (RMON) to gather these statistics. The GSD must tightly restrict the systems able to poll using these methods because of the inherent insecurities of the aforementioned monitoring protocols. By trending these metrics, it may be possible to determine when it is time to increase bandwidth or purchase systems that are capable of meeting the new throughput or processing needs.

26.5.1.2 Availability. When GSDs are unavailable, the network functionality can dramatically diminish. A simple test of system availability is using ICMP to "ping" one or more interfaces to ensure the device itself is responding. However, this approach can be deceptive. Just because the device itself responds, does not mean it is properly forwarding traffic. It is also advisable to send probes (e.g., ICMP, traceroute, or other queries) to something on the other side of each interface to ensure the other device is actually receiving the packets to ensure valid results. This approach provides a better overall picture of the GSD availability.

26.5.1.3 Integrity. The ability to trust network security systems components is vital. The possibility of a root kit compromising a firewall or GSD is now a reality. These systems must have the ability to protect against modification of system components such as ceasing operation and/or alerting the change. If this embedded functionality were unavailable, it is possible to write a script to generate cryptographic hashes of critical system components and verify against a known trusted version.

26.5.2 Policy. The GSD policy is the core definition for providing and protecting allowed paths. Most security systems process packets starting at the beginning of the policy and continue until there is either a match or reaching the end of the rule base (which should be an explicit "deny any"). As discussed later in this chapter, there are situations where rules may process before or after the main rule base.

Centralized management consoles provide intuitive GUIs to configure and easily manage one or more firewall and GSD policies. Certain platforms also provide the ability to manage policies directly from the device.

26.5.2.1 Defining Allowed Paths. Allowed paths identify specific protocols used to implement communication. In a typical Internet environment, business services require allowed paths such as HTTP(S), SMTP, and DNS. These requirements will vary, but for an environment, each allowed path should directly relate to the required service.

Starting from an implicit or explicit (depending on the platform) "deny any" rule, allowed paths will be added as "allow" rules, such as PERMIT HTTP, with specifics determined by the following sections.

Although network addressing does not provide effective authentication of systems or users, restrictive endpoints can make it much more difficult for an attacker to exploit an otherwise straightforward vulnerability. It is also important to identify the endpoints carefully, particularly in cases where these endpoints might reside on internal rather than extranet or utility zones.

The direction of traffic, indicated by the source of the connection initiation, is useful for the rule definitions for several reasons. First, rules can be written so that only responses to internally originated allowed paths are allowed in from the untrusted zone, rather than explicitly permitting the protocol bidirectionally. In addition, the firewall may process rules at different times based on design and/or configuration.

26.5.2.2 Complexity of GSD Policies. Standard firewall rules operate on simple Boolean principles. For example, allow or deny network traffic that is going from host or network X to Y on port Z. The complexities required of GSDs evaluating network traffic are dramatically higher. For example, this evaluation could be a combination of a Boolean test to verify an inbound email address is from a trusted source, verify message contents are acceptable, and scan an attachment for viruses. Administrators must understand the higher-level protocols to ensure that the GSD policy matches the types of protections expected and required. As the number and complexity of the rules increases, so do the processing requirements for the GSD.

Beyond the basic firewall capabilities, GSD policies typically include per-rule enforcement of the additional security measures. For example, the GSD only uses network and/or transport layer filtering to restrict access to the organization's VPN terminators. The policy may include WAF protections, but only for the utility zones. Inbound and outbound Web and file transfer traffic may have the additional requirement for NIPS inspection. Although this level of customization adds some administrative overhead, selectively (rather than broadly) applying additional protections in this manner can reduce the performance impact of the added protections.

26.5.2.3 Change Management. Whether managing one or one hundred policies, an essential element is to have a process to track policy changes. Change management can be cumbersome but has several advantages. First, this provides back-out information if a change were to cause issues. Second, it provides an audit trail of

who requested the change (and why) in addition to who changed the policy. Lastly, it provides a method for streamlining change requests by providing a single method for accepting, reviewing, implementing, and validating policy changes.

26.5.2.4 Secondary Evaluation. Making policy changes tends to be a simple process, such as adding HTTPS access for a new extranet server. However, the more complicated the change, the more likely for mistakes to occur. Having another administrator evaluate the proposed change brings a fresh perspective that may catch a small discrepancy that could have negative consequences. Another step in the change management process is validation across multiple operational groups—providing an opportunity to detect potential conflicts with other parts of the infrastructure due to the change.

26.5.2.5 Auditing/Testing. Is your firewall or GSD working as expected? Is the GSD catching the most recent malware-infected attachment? Would you bet your organization's future on it? These questions are simple yet powerful reminders that you must verify that the network-security devices are working as expected. In addition to on-going and in-depth log reviews, a proper regiment of auditing and testing will help to answer these questions. In addition, management consoles should provide an audit log that tracks who makes changes to any part of the environment.

Vulnerability assessment and penetration testing are effective methods to determine the validity of policy rules. For example, the firewall should not allow an HTTP connection to a system with only SSL/TLS access configured. If the HTTP access succeeds, is the failure on the tested rule, is there another rule higher in the rule base erroneously allowing this access, or is something unexpected occurring? The tools and processes must be in place to answer these questions before someone else does.

26.6 MAINTENANCE. Patching workstations and updating malware definitions/signatures are common practices, but this process is even more crucial for systems protecting the network. It is essential to test each update's validity and stability to ensure faulty or rogue changes do not interrupt operations.

26.6.1 Patching. No system is inherently and infinitely secure, and it may be possible to subvert a security device due to system vulnerability. It is crucial to monitor GSD Websites for the most recent operating system and system component patches. By also monitoring vulnerability and exploitation information sources, it may be possible to ascertain more timely information. This additional research could result in being able to determine and implement a temporary solution until the vendor can provide a permanent fix.

Vendors can take weeks and months to develop and release a patch. In certain cases, third-party patches will become available for un-patched, actively exploited vulnerabilities. Although using these is an option, it is inadvisable due to the inability to verify the patch's integrity and safety.

26.6.2 Pattern Updates. Threats evolve at will. Automatic updates provide the smallest delta of exposure for pattern-based malware signatures. However, blindly trusting these updates in a production environment may have adverse effects. To avoid issues, establish a procedure that enables new signatures in monitor only mode to ensure there are no adverse effects before implementing full protection. If automatic updates

are not a viable option, then testing requires a lab environment or use of noncritical systems to vet the new patterns/signatures.

26.6.3 Logging and Alerting. Whether remaining local or transferring to a centralized management system, logs take up a large amount of disk space, and if left unviewed, are worthless to keep. There must be a log review process. Instead of trying to determine the anomalies from an entire log set, it is helpful remove known traffic to help expose the unknown and potentially malicious. For example, if you only expect outbound SMTP traffic from a single network, you would filter out those known items and more easily see if the other unauthorized hosts or networks are originating SMTP traffic.

Organizations with a large Web-based footprint will generate a large number of HTTP/HTTPS-related events because each new connection creates a log entry. By having a GSD do protocol inspection and logging, there is the potential of gaining a better understanding whether packets are expected or malicious. If there are other systems tracking and/or correlating these connections, it may be useful to turn off logging on selected GSD rules to reduce the overabundance of logs that are unnecessary to review.

A security incident and event management (SIEM) system provides an additional method for collection, aggregation, and consolidation of logs from many types of devices. The SIEM leverages baselining and configurable rules to correlate the logs and provide real-time incident-based alerting.

Alerting is complementary to and usually depends on the logging mechanism. Once the firewall or GSD generates a log entry, the administrator can configure alerting options when certain log conditions or threshold changes exist. For example, the administrator can configure an alert to send a notification if an internal system process changes state. By carefully determining alerting thresholds and notification methods, there is less opportunity to burden the administrator(s) with nuisance alerts.

26.6.4 Secure Baseline Configuration. Another crucial step to protecting network-security devices is to create a secure baseline configuration. In certain cases, the GSD vendor provides a hardened and/or proprietary version of an operating system. This should not be a "green light" that the device is secure and ready for production. Rather, this provides even more reason to take the time to evaluate the security posture of the device thoroughly.

Once a secure configuration is set, this then becomes the baseline configuration for the remainder of the systems. This provides a standard configuration helping to ensure each device functions and secures in the same manner. There should also be a process to ensure the integrity of the secure configuration as well as verifying and testing when valid updates occur.

26.6.5 Default Configurations. Default system and policy configurations vary widely; some implicitly deny, while others may implicitly allow. In no case should any network-security device deploy into production with a default configuration. Ensure that you determine whether the network-security device fails in an insecure or secure posture. Inline devices may fail insecure in an attempt to maintain traffic flow in the case of failure. Depending on the network, this may or may not be the optimal response. Although potentially disruptive, a fail insecure posture reduces the likelihood of anomalous or malicious traffic passing undetected.

Implied Rules. Implied rules are separate from but are part of the default configuration. The network-security device will process the implied rules prior to packets getting to the policy rules. The implied rules may allow the GSD to process specific and known administrative traffic without using processing cycles to go through the entire rule base. It is essential to determine whether these rules exist as well as whether they match the desired security posture. If not, disable or modify the implied rules to fit specific protection needs.

26.6.6 Harden the Device. Protecting the network-security device is a critical part of the security of the overall infrastructure. The two most common ways to reduce these exposures are through the administrative console and validation using vulnerability assessment. The management or system console can provide information such as default listening ports, services running, and the like. If a service is not critical to the functioning of the firewall, disable it to remove any threat of it becoming a compromise vector. Even with the service disabled, ensure you continue to apply patches for those services. Only specific administrative hosts should have direct access to the system. As with the GSD policies, it is useful to validate that configuration changes are actually making the device more secure. By conducting a vulnerability assessment, you are able to determine whether unexpected services are still available or additional vulnerabilities exist.

26.7 DISASTER RECOVERY. The impact of a GSD outage can range from an annoyance to critical event depending on implementation, necessity, and disaster recovery planning. Since these systems represent part of the backbone of the security infrastructure, it is a necessity to provide continual, protected network access.

26.7.1 Fail-Over/High Availability. High availability (HA) architectures are essential in situations where reliable access is a necessity. The HA deployment will typically have an Active/Standby pair; where the active member is processing live packets and the standby member is ready to take over in the case of a failure of the active member. If the standby member is continually synchronizing connection information from the active member, there would be little to no loss of connectivity if the active member fails because the standby would then become active and already have the connection information to allow continued processing of the current sessions and service new sessions.

26.7.2 Load-Balancing. Distributing the load between multiple systems is another way to reduce an availability exposure. If the load equally distributes between two or more systems, the failure of one does not eliminate all access. One caveat is that the load balancer must continue to route each connection to the same system on the back-end to maintain connection information. Otherwise, it may be possible for the load balancer to route traffic to another system that does not know about the established connection (e.g., asymmetric routing).

26.7.3 Backup/Restore. Backup and restore functionality is the most rudimentary disaster recovery method. If the system crashes or suffers a compromise, there must be a way to recover the system in an efficient manner. By having a reliable backup process, it becomes easier to restore the device configuration in the event of a failure. Depending on the vendor, backups may consist of a text file or something more complex, like a GZIP file containing the crucial configuration information. The backup

process must include a method to move the backup to another system to eliminate a single point of failure. The restore process may involve restoring the operating system and/or the security device configuration. This may be as simple as uploading the backup file or may require an out-of-band connection to do prework before uploading a preserved configuration.

26.8 NETWORK-SECURITY DEVICE EVALUATION. As digital threats continue to evolve, so must network-security devices. Internet connectivity costs continue to decline and so does the ability to detect complex threats. It is a mistake to think the current security systems and infrastructure will be viable for an extended length of time. In essence, the infrastructure is continually under review. One crucial time to revisit these items is when it comes time to replace security devices.

Effective security lies at the intersection of protection/functionality, usability, and cost. Every situation has unique challenges and opportunities—the goal is to develop the best security solution to meet the organization's ongoing risk requirements. The following sections provide a framework for evaluating network-security devices currently in-service or during a Request for Information (RFI). Please note that this section addresses the process over the lifecycle of the GSD architecture instead of providing detailed guidance on security-inspection capabilities. These guidelines provide a good basis for thoughtful analysis of site-specific requirements.

26.8.1 Current Infrastructure Limitations. It may be technically possible to do a one-for-one replacement of an existing firewall with a GSD, but it is important to review the current infrastructure to ensure maximum effectiveness of the new device(s).

- Are the current network-security devices past useful life?
  - This includes devices that are out of warranty, no longer vendor supported, and unable to meet current requirements.
- Is the Internet-facing architecture limiting GSD deployment options?
  - Having only a few publically routable IP addresses can limit the number of external service zones without extensive use of NAT.
  - Placing the GSD in an allowed path of all encrypted traffic may relegate the GSD to being more packet filter than extensible security solution.

26.8.2 New Infrastructure Requirements. The dynamics of network and security infrastructures shift constantly. Trends of rapidly escalating bandwidth usage or an increased reliance on encrypted communications are common. The prospect of adding the GSD capabilities have a dramatic influence on these decisions.

- Are there new demands for increased bandwidth, encryption, or application layer enforcement?
  - Determining new demands include observed or anticipated/projected growth.
- Is Internet-access redundancy becoming a necessity?
  - It is important to consider a GSD with built-in redundancy capabilities, such as Active/Passive failover or clustering.
  - Is your infrastructure able to support diverse carriers and/or routes to the Internet?

26.8.3 Performance. Technological obsolesce continues to haunt organizations attempting to grow to meet internal and external demands.

- Does the existing device have excess capacity for future needs?
  - Security devices running at high processor and/or memory usage are good indicators that the system is at or reaching its maximum performance threshold.
  - Without extra processing capacity, a modest increase in network traffic or enabling a new feature could cripple an already resource starved device.

26.8.4 Throughput. The need for bandwidth is relentless. Bandwidth of 10 Gigabit is no longer sufficient in the core since the price and availability of 40 and 100 Gigabit is already increasingly obtainable. Standard 802.11ac is ready to eclipse the establish 802.11n wireless standard and offer claimed speeds of a gigabit per second (or greater), leading organizations to questions their use of wired connectivity. Organizations should take a realistic approach before paying for higher bandwidth solutions.

- Does the device meet bandwidth requirements?
  - The GSD must be able to sustain the necessary peak traffic loads without interruption or service degradation.
  - Are aggregated network interfaces supported to support modest increased bandwidth rather than paying for higher speed interfaces (e.g., multiple Gig interfaces versus 1×10 Gig)?
- Does the device have excess capacity for future needs?
  - Modular devices are able to increase memory, add network interfaces, or increase bandwidth by adding or replacing a hardware module.
  - Peaks in network traffic are inevitable. Having additional bandwidth capacity for existing GSD network segments will reduce the impact of these situations.

26.8.5 Implications of Added Features. Many vendors tout throughput at (or near) wire speeds. However, these statistics do not necessarily mirror reality as the standard throughput tests always consider the varying packet sizes and protocol mixtures. In actuality, packet types and sizes vary constantly; be sure to review vendor claims through independent testing organization and, if possible, internal testing.

- What performance impact does activating additional security features create?
  - The greater the number of additional functions—typically resulting in more detailed payload inspection—the increased negative impact on the device's processing power and bandwidth.
  - If requiring encryption, hardware acceleration, if available, will reduce much of this specific performance impact.
  - Be sure to evaluate the vendor's and independent documentation showing measured performance against varied traffic flows such as IMIX (Internet Mix) and EMIX (Enterprise Mix).

- What level of additional administrative overhead will the features create?
  - Consolidating management of multiple security functions does not necessarily make them easier to manage. This also adds complexity depending how each subsystem integrates.

26.8.6 Management. GSDs typically require substantial planning, configuration, monitoring, and maintenance. Having a robust management platform in distributed GSD environments is essential to the success of the implementation.

- Does the current environment have the management features needed or required?
  - Self-evaluation of the current management environment is fairly easy, as there are already established likes, dislikes, nice-to-haves, must-haves, and must-nots.
- Is there the need/want for distributed or centralized management?
  - The size and type of GSD deployment is important to consider. It probably does not make sense to use an extensive management platform for a few devices. With larger geographically distributed environments, having centralized control of all devices is necessary for efficient and effective command and control.
- Are redundant management systems needed?
  - Redundant management systems ensure administrative control and logging of GSD devices in the case of a failure. Be sure to review management systems redundancy options to ensure the secondary system has functional parity.
- Is the ability to manage the system from both a centralized management console and/or directly from the device a necessity?
  - Certain types of GSDs allow system and policy changes to be made directly from the device. This feature can be useful during testing or troubleshooting but may cause confusion or misconfiguration if not properly understood and managed.
- Is encryption a requirement to protect management, policy, and logging functions?
  - The need to encrypt sensitive information is an important consideration for transmission of GSD management and logging functions. If sent clear-text, it is possible for someone to monitor these communications and gain significant intelligence about the network security architecture.
- Is detailed reporting mechanism a necessity?
  - Most logging systems provide the ability to filter data, but it may be useful to do more in-depth reporting. Some example reports include rule base hit percentages (e.g., which rules receive the highest percentage of hits) or traffic usage statistics (e.g., percent usage by protocol, network segment, host, and/or user).
- What level of granularity is necessary for management permissions?
  - With only a few administrators, this may not be as significant, but larger deployments may have distributed administrators. The GSD management system must provide role-based security with the ability to restrict access to certain tasks, areas, and/or logs. It may be useful to compartmentalize the systems and allow other support teams view-only access for troubleshooting and auditing purposes.

- Are both GUI and/or command-line policy or device management options required?
  - Certain GSD brands have the option of configuring devices using a more traditional command line interface and can be useful for quick troubleshooting and scripted changes. A drag-and-drop graphical environment uses the central management systems to easily create, change, and view policies and logs. If both options are available, ensure both options have feature parity.

26.8.7 Logging and Alerting. Although the GSD's primary responsibility is denying malicious traffic, logging and alerting play a critical role in the overall security posture. Logging is essential to investigating incidents, viewing trends, and may even be a regulatory requirement.

- How do you want to be able to filter logs?
  - Collecting, searching, and reviewing logs in a text-based format has advantages, including consolidation and scripting for log reviews. Text-based log require scripts to parse through logs but do have the advantage of being able to use complex queries and the ability to analyze logs on virtually any system.
  - Graphical user interfaces (GUIs) provide the ability to filter logs based on one or more criteria but are weak when attempting complex or recursive queries. A GUI increases the aesthetic appeal of the logs and can ease viewing by grouping or filtering different types of events, as well as adding color and/or graphics that help catch the eyes' attention more quickly. Although useful, the GUI may be slower to load and may restrict the ability to create complex queries.
- Do you want compressible logs?
  - Even with well-maintained logging standards and grooming processes, the sheer volume of GSD logs alone can be overwhelming. Text-based logs typically provide high compression ratios and easy archival. Some of the proprietary formats may automatically provide some compression but may not be able to match that of nonproprietary formats.
- How do you want to store logs?
  - GSD logging options include storage on the local device, on an integrated logging and management platform, or on a separate storage device. Local storage creates the risk of data loss if the GSD fails or potential for system impact if the logs fill the local storage. Sending the logs to other systems will increase network traffic, but this provides easier log management over the information's lifecycle.
- What types and granularity of alerts do you need?
  - Alerting helps to bring information to life without the need to watch logs every minute of the day. It may be possible to create email or SMS alerts an event occurs breaches a configured threshold (e.g., the GSD network throughput is more than 85 percent of the maximum possible or a host is attempting to attack the GSD).
  - It may be possible to develop detailed alerting conditions (e.g., criteria dependencies) to minimize erroneous alerts.

- Do you have a legal and/or regulator requirement to keep logs for a set period of time?
  - Internal and/or regulatory requirements may make it necessary to keep this data ranging from months to many years.

26.8.8 Auditability. Audit records provide an accounting of what changes occurred, when, by whom, and so forth. Having an integrated audit trail of GSD environment changes can be invaluable during an incident.

- Do you want to be able to audit device changes?
  - Changes to the device itself do not occur nearly as often as changes to policies. Having the ability to collect information about these changes can help uncover system instability, accidental misconfiguration, and/or malicious intent.
- Do you want to be able to audit policy changes?
  - GSD policies tend to change rapidly as new access requirements develop, new capabilities or requirements appear, and the ever-changing threat landscape progresses. Policy audit trails help to resolve mistakes and provide archival information of what changes occurred and when.
- What level of audit granularly is necessary?
  - Data such as date, time, type of change, to what system(s), and from what administrator are all valuable pieces of information. Increasing auditing granularity can increase the overall log volume and processing overhead.
- Should the audit logs be independent from the firewall event logs?
  - Having a logging system that automatically separates audit logs from other event logs helps ease viewing and searching. This also creates the ability to create more granular permission levels (e.g., who can see which type of logs).

26.8.9 Disaster Recovery. With GSDs providing consolidated security functions, incorporating disaster recovery (DR) options into the GSD environment helps ensure it does not become a single point of failure for the entire organization.

- What are your specific disaster recover requirements?
  - As mentioned earlier, different types of redundancy help during DR scenarios. High availability setups such as active/passive failover and clustering can help to increase overall availability.
- Does the device have easy-to-use backup and restore mechanisms?
  - The backup mechanism should provide a simple interface to complete backup and restores as well as transfer them off the local device for storage and archival.
- Do you want the ability to do automated restores?
  - Certain GSDs have the ability to do an automated or semi-automated restore using scripts or through pushable configurations from the management system.
- Should the system fail secure or insecure?
  - It is important to determine if the uptime advantage of a fail insecure system outweighs the risk of network traffic passing uninspected or the downtime associated with a fail secure system.

- Do you need out-of-band access?
  - Having an out-of-band access solution such as a direct serial connection and/or dial-in solution may provide the ability to monitor and/or change the system during the network outage.

26.8.10 Usability. Having a management environment that has every feature imaginable is worthless if it is too complex and convoluted to operate and maintain.

- Is the management console intuitive?
  - The management console should have a consistent and simple interface to manage the devices and policies. The more intuitive the interface, the less time the administrators will need to spend figuring out how to complete otherwise easy tasks or troubleshooting issues.
- Are the primary functions easy to accomplish?
  - The interface should make it easy to follow and edit policies as well as view and filter logs.

26.8.11 Learning Curve. By choosing a GSD vendor with an intuitive interface and functionality, it should be much easier to learn the environment.

- Is training required to learn the new device(s) and management platform?
  - Even with the best interface and features, it may be necessary for the administrator(s) to get training to use the system effectively. Check with the vendor to see what training (basic through advanced) is available and at what cost. If possible, determine the effectiveness of the instructor(s) through referrals.
  - If some of the security features are new to the organization, expect a steeper initial and increased ongoing learning curve due to the breadth and depth for each additional protection measure.
- How does the vendor approach security architecture?
  - Every GSD vendor has a little different view and implementation of security features. Some focus primarily on allowing or denying traffic, although others treat network traffic as flows between zones. Moving from one vendor type to another can be confusing and could slow or confuse the conversion.

26.8.12 Features. Although the GSD platform may be able to provide a single source for network-security protections, it is important to remember that it may not be cost or resource effective to meet the organization's security needs and posture.

- Are you looking for a device to be an all-in-one solution?
  - The ability to provide a single solution security platform has merits, but each additional feature has an impact on performance and availability, and when implemented improperly, increases the potential to cause issues.
- What features do you want the GSD to have?
  - Common UTM features may include firewall, IPS, VPN, anti-malware, anti-spam, DLP, and Quality of Service (QoS).
  - Common NGFW features include port-independent application identification and enforcement for sub-features of a particular protocol.
  - Does the platform support full IPv6 enforcement across all security functions?

26.8.13 Price. Be sure to understand the cost implications of the entire GSD infrastructure over its lifetime.

- What will be the *total* cost of ownership (TCO) of the upgrade or replacement?
  - Areas for considerations include: hardware purchase price; price for requested security features (additional features, if added later, will increase the overall costs and may require new hardware); management infrastructure (hardware and software); hardware and software maintenance fees; training costs; direct or third-party support contract; potential conversion downtime; and learning curve.

Initial Cost

- How much will the new GSD equipment cost?
  - This includes the price of the GSD, expansion cards, and management devices.
  - An appliance may be a more cost-effective alternative long term if there are multiple deployment options as the vendor is solely responsible for issues.
- How much extra will shipping cost?
  - If this is a distributed deployment, there may be extra shipping costs and tariffs associated with getting the equipment to international locations.
- How much will it cost to purchase the features I need?
  - This is highly dependent on the GSD. Certain GSDs come with everything out of the box, while others provide some basic functionality including the firewall, VPN, and basic protocol inspection but charge extra to use the full functionality of the existing items and/or to add new functionality.
- How much is it necessary to invest in training to learn the new solution?
  - Training often includes attending offsite courses (which may require additional travel expenses). In some cases, it may be more cost effective to bring a trainer in-house to provide a more tailored training program. A train-the-trainer methodology can be a money-saving option, depending on team size and funding.

Ongoing Costs

- How much is the yearly hardware and/or software maintenance?
  - The maintenance contract may be negotiable based on length of term and costs (potentially based on a percentage of the purchase).
- What level of service do you require?
  - GSD vendors provide different service levels to correspond with response times, escalation ability, and access to additional information. In accordance, the higher level of service, the more this will cost.
- Can the provider support you at each of your locations?
  - If necessary, major issues may require a vendor or vendor partner to come onsite to complete a repair. If available, determine the cost (time and travel) before using this type of service.
  - If you have international deployments, be sure to understand the vendor's ability to provide general warranty repair and/or onsite service for those locations.

- Is there a requirement for an onsite spare?
  - If the DR plan requires quick recovery, having an onsite spare may help. The disadvantage to having a spare (not a redundant pair) is that it most likely will become an idle asset. Be sure to weigh this versus the cost of a higher-level service contract or local availability of replacement parts.
- Do you need additional consulting/support time?
  - Vendors and vendor partners may provide blocks of time where a dedicated resource can come onsite for additional support, planning, and platform reviews.

26.8.14 Vendor Considerations. Being able to trust that a vendor's product will protect the organization is crucial. Be sure to feel comfortable with every aspect of the product and service before turning over the network security reins. In addition to getting the features required, the vendor providing the solution should have a solid security foundation, financial stability, and adequate support resources.

- Will this vendor meet the organization's current and future needs?
  - This key question focuses on the vendor's ability to provide a quality product, support that product, and grow with your needs. Be sure to evaluate the vendor's product roadmap to see how the product (hardware, functionality, cost, etc.) will evolve and when currently unavailable features will be available in the product.
- Does the vendor already support or have experience with deployments similar to yours?
  - If the vendor already has customers with an infrastructure and needs similar to yours, this may provide additional ideas and/or information to make the deployment more successful.

26.8.14.1 Reputation. Although infighting and opinions about the best security platforms abound, the selected vendor should be an active member of the security community.

- Does this vendor have a good reputation in the security community?
  - There are no shortage of product reviews and comparisons to help determine whether a product or vendor has a solid reputation for quality and service. These are not the ultimate arbiter of a good product but are a foundation for additional research.
- Will they provide references?
  - Check with the vendor to see if another company (or two) will provide a reference as to their experience with the product and service.
- What is the vulnerability history of the current and previous systems?
  - Check vulnerability and exploit monitoring sites to determine vulnerability history.
- What is the mean time to correct vulnerabilities?
  - Although the optimal time to resolve the vulnerability is immediately (or prevent it before it happens), check to see how long it takes vendors to respond to and resolve vulnerabilities. This includes verifying the fix corrected the problem the first time.

- What is the mean time to correct application identity issues and/or gaps?
  - The complexity and changing nature of application layer protocols requires active vendor participation. Determine the process, timelines, and potential costs associated with the vendor correcting (or supporting) application identity needs.

26.8.14.2 Support Options. Due to the complexity and breadth of functionality these devices possess, the selected vendors must have a robust support infrastructure to resolve issues.

- How experienced are the front-line and higher-tier technicians?
  - Ask the vendor to provide information regarding training and experience of the different support technician levels.
- Are there different tiers of initial call support?
  - Determine whether all support calls originate through the main support/service desk or if it is possible to have a dedicated (typically shared among other organizations supported) resource.
- Do you have the ability to escalate-on-demand?
  - Determine whether it is possible to escalate an issue immediately (and at what cost) if the local administrators complete standard troubleshooting steps.
- Are you willing to work with systems during beta testing?
  - Depending on risk tolerance, being a beta partner with the vendor may provide early access to fixes and enhancements.
- What are current and previous clients' experiences with support quality and availability?
  - Again, ask for references to get real-world perspective.

26.8.15 Managed Security Service Providers. Organizations have the ability to transfer varying levels of internal network security responsibilities to a Managed Security Service Provider (MSSP). Use of the MSSP may be an opportunity to supplement off-shift log reviews, consolidation, and/or alerting. In addition, there may be an opportunity to transfer maintenance and change control to alleviate internal resource constraints and/or knowledge gaps. Choosing this type of service requires detailed investigation.

- What is the MSSP's reputation/experience/workload?
  - The MSSP should have trained and experienced personnel to support your GSD infrastructure.
  - Determine the number of clients per support resource and if other MSSP locations can continue support and operations if a failure occurs at one MSSP location.
- Is it possible to do this securely?
  - Ensure the MSSP complies with secure storage of collected information and access to the GSD management infrastructure.

- If necessary, how will the MSSP follow change control?
  - Determine the change control process before turning over operational responsibility. This includes determining change reviews, change windows, change approvals, service-level agreements (SLAs) regarding time to complete, etc.
- Can the MSSP meet your SLA requirements?
  - Set clear expectations during contract negotiations when establishing SLAs. The more stringent the SLA, the more likely the cost is to increase. Develop metrics to assess the MSSP's effectiveness in achieving the SLA requirements.
- At what cost?
  - The cost of doing business with the MSSP increases based on the number of services and/or devices. Be sure to investigate the contract to understand all potential fees or additional requirements that may drive up the cost after the signing of the initial contract.

26.9 CONCLUDING REMARKS. The requirements for network security are changing constantly. The organizations are under continued pressure to meet customer and business partner demands for information and access. Technologies such as mobile devices and extranets continue to blur the previous accepted reality of true network borders. Perimeter security is no longer as easy as installing a firewall or requiring proxy services for outbound connections. Reliance on a specific technology or security ideology is insufficient to protect the information and systems integral to the success of an organization. The frequency and veracity of both external and internal attacks requires a more robust and flexible mechanism to combat these increasing threats, and the GSD leapt forward to provide the current generation of perimeter protection technologies.

The GSD retains all of the previous firewall-based functionality such as allowed-path control, VPN services, and network address translation while providing the flexibility to integrate additional services designed to provide visibility into and protection from a multitude of threats. The addition of anti-malware capabilities provides a new layer of protection by relieving hosts from being the only location providing full detection, prevention, and remediation services. The integration of proxy services, application and content control, and intrusion prevention allows GSD deployments to simplify the network-security architecture and opens up new inspection and protection opportunities.

Today's threats provide little insight into how protection measures will need to evolve to meet the next-generation attacks and attackers. To remain a viable security option, the GSD security vendors must remain agile. The integration and implementation of new protection measures will need to be simple and seamless. As processing power continues to increase, so will the capabilities of the GSD to take on greater workloads and complexity. Basic content inspection will give way to a greater understanding of and protection for data context, value, and flows. The support and protection of worldwide networks will require providing the full existing set of GSD functionality for IPv6 traffic.

No matter the threat, no matter the network, no single device is a silver bullet to complete security. Each organization must evaluate all avenues of protection to ensure that the technologies deployed meet the security functionality required.

26.10 FURTHER READING

Forouzan, Behrouz A. *TCP/IP Protocol Suite*, 4th ed. McGraw-Hill, 2009.

McClure, Stuart, Joel Scambray, and George Kurtz. *Hacking Exposed*, 7th ed. McGraw-Hill Osborne Media, 2012.

Scarfone, Karen and Paul Hoffman. "Guidelines on Firewalls and Firewall Policy: Recommendations of the National Institute of Standards and Technology." NIST Special Publication 800-41 Revision 1, 2009. http://csrc.nist.gov/publications/nistpubs/800-41-Rev1/sp800-41-rev1.pdf

Thomson, S., T. Narten, and T. Jinmei. "RFC 486: IPv6 Stateless Address Autoconfiguration." Internet Engineering Task Force, 2007. https://tools.ietf.org/html/rfc4862

Narten, T., E. Nordmark, W. Simpson, and H. Soliman. "Neighbor Discovery for IP version 6 (IPv6)." Internet Engineering Task Force, 2007. http://tools.ietf.org/html/rfc4861

Deering, S., W. Fenner, and B. Haberman. "Multicast Listener Discovery (MLD) for IPv6." Internet Engineering Task Force, 1999. http://tools.ietf.org/html/rfc2710

Johnson, D., C. Perkins, and J. A. Ericsson. "Mobility Support in IPv6." Internet Engineering Task Force, 2004. http://tools.ietf.org/html/rfc3775

CHAPTER 27

INTRUSION DETECTION AND INTRUSION PREVENTION DEVICES

Rebecca Gurley Bace

27.1 SECURITY BEHIND THE FIREWALL 27·2
- 27.1.1 What Is Intrusion Detection? 27·2
- 27.1.2 What Is Intrusion Prevention? 27·3
- 27.1.3 Where Do Intrusion Detection and Intrusion Prevention Fit in Security Management? 27·3
- 27.1.4 Brief History of Intrusion Detection 27·4

27.2 MAIN CONCEPTS 27·4
- 27.2.1 Process Structure 27·5
- 27.2.2 Monitoring Approach 27·5
- 27.2.3 Intrusion Detection Architecture 27·5
- 27.2.4 Monitoring Frequency 27·5
- 27.2.5 Analysis Strategy 27·6

27.3 INTRUSION PREVENTION 27·6
- 27.3.1 Intrusion Prevention System Architecture 27·6
- 27.3.2 Intrusion Prevention Analysis Strategy 27·6

27.4 INFORMATION SOURCES 27·6
- 27.4.1 Network Monitoring 27·7
- 27.4.2 Operating System Monitoring 27·7
- 27.4.3 Application Monitoring 27·7
- 27.4.4 Other Types of Monitoring 27·7
- 27.4.5 Issues in Information Sources 27·8

27.5 ANALYSIS SCHEMES 27·8
- 27.5.1 Misuse Detection 27·8
- 27.5.2 Anomaly Detection 27·9
- 27.5.3 Hybrid Approaches 27·10
- 27.5.4 Issues in Analysis 27·10

27.6 RESPONSE 27·10
- 27.6.1 Passive Responses 27·10
- 27.6.2 Active Responses: Man-in-the-Loop and Autonomous 27·11
- 27.6.3 Automated Response Goals 27·11
- 27.6.4 Investigative Support 27·12
- 27.6.5 Issues in Responses 27·12

27.7 NEEDS ASSESSMENT AND PRODUCT SELECTION 27·13
- 27.7.1 Matching Needs to Features 27·13
- 27.7.2 Specific Scenarios 27·14
- 27.7.3 Integrating IDS Products with Your Security Infrastructure 27·14

27.7.4 Deployment of IDS Products 27·15
27.8 CONCLUSION 27·16
27.9 FURTHER READING 27·16
27.10 NOTES 27·17

27.1 SECURITY BEHIND THE FIREWALL. Even today, when asked how they would go about securing a computer or computer network, most people mention firewalls, the first widely accepted network security devices. As network security has become a nonoptional facet of system management, firewall mechanisms of various sorts have become a standard fixture in many networks.

As in any complex protection function, firewalls are necessary but not sufficient to completely protect enterprises from security breaches. Expecting firewalls to provide complete protection is tantamount to expecting guards at the gates of corporate campuses to prevent destructive behavior on the part of those operating vehicles within the facility. A lot can happen as traffic flows between hosts on internal networks; even items that appear benign to gatekeepers can be subverted to carry destructive payloads. Thus, most modern network security architectures include intrusion detection and intrusion prevention systems.

Intrusion detection systems (IDSs) are software or hardware systems that automate the monitoring of events occurring within a computer system or network. IDSs not only collect and synchronize records of these events; they also analyze them for signs of security violations. In strictest terms, vulnerability assessment systems (VASs) are a special class of IDSs in which the system relies on static inspection and attack reenactment to gauge a target system's exposure to specific security vulnerabilities. Intrusion prevention systems (IPSs) are another special class of IDSs in which the system is designed to react to certain detected attacks in a prespecified way.

In recent years, the realm of intrusion detection has broadened and deepened, driven by a number of factors. Security personnel, architectural and operational alike, have gained more experience with IDS technologies, working with commercial product providers to expand product capabilities and management schemes to fit current needs. As operational personnel have become more comfortable in using these systems, they have moved certain IDS capabilities to core network management venues. Finally, threats have evolved, driving needs that IDSs are uniquely qualified to address.

The evolution of IDS has resulted in some changes in nomenclature and tradecraft associated with it. One of these changes is that vulnerability assessment is considered a stand-alone discipline, driven by market needs that are often different from those influencing IDS. Another is that intrusion prevention has evolved as a stand-alone product category, offering the ability to respond automatically to certain classes of attacks. Thus, vulnerability assessment is treated as a separate topic in this *Handbook* (covered in Chapter 46), while both IDS and IPS are discussed in this chapter.

27.1.1 What Is Intrusion Detection? Intrusion detection is the process of collecting information about events occurring in a computer system or network and analyzing them for signs of *intrusions*. *Intrusions* are defined as violations of security policy, usually characterized as attempts to affect the confidentiality, integrity, or availability of a computer or network. These violations can come from attackers accessing systems from the Internet or from authorized users of the systems who attempt to overstep their legitimate authorization levels or who use their legitimate access to the system to conduct unauthorized activity.

Intrusion detection systems are software or hardware products that automate this monitoring and analysis process.

27.1.2 What Is Intrusion Prevention? Intrusion prevention is the process of coupling intrusion detection (as defined) with specified responses to certain detected intrusion scenarios. The triggering events can be viewed as a special subset of intrusions, and are often characterized in richer quantitative and qualitative terms than more generic IDS triggers. For example, a specific IPS might focus on monitoring certain types of network traffic. When the rate of a particular traffic type exceeds the anticipated threshold, the IPS would react in a prespecified way (e.g., limiting the rate of subsequent traffic of that type).

27.1.3 Where Do Intrusion Detection and Intrusion Prevention Fit in Security Management? Intrusion detection is a necessary function in most system security strategies. It (and its offshoot, vulnerability assessment) is the primary security technology that supports the goal of *auditability*. "Auditability" is defined as the ability to independently review and examine system records and activities to:

- Determine the adequacy of system controls
- Ensure compliance with established security policy and operational procedures
- Detect breaches in security
- Recommend any indicated changes[1]

The presence of a strong audit function, in turn, enables and supports several vital security management functions, such as incident handling and system recovery. Intrusion detection also allows security managers a flexible means to accommodate user needs while retaining the ability to protect systems from certain types of threats.

There is some debate over whether IPS will displace IDS in the security management lineup. This displacement is unlikely for the time being because of the binding between IDS and audit. As security operations become more tightly integrated with traditional system administration and operations, audit functions are necessary to support root cause analysis in diagnosing and addressing system failure. Badly tuned security devices can create problems in operations; it is important to be able to identify and correct such issues quickly and appropriately. Audit functions are also necessary to measure the effectiveness of security measures in mitigating security threats. Although it may appear that transparently detecting and blocking attacks is the optimal security process, such transparency—that is, lack of an audit trail—interferes with demonstrating effectiveness of security measures against real threats. In other words, in order for security management to justify a budget for security measures, it must be able to document and quantify the suitability and effectiveness of such measures. Therefore, regardless of the effectiveness of IPS in blocking attacks, IDS will likely always be necessary to support these audit functions. IDS provides the baseline information on the attack profiles and frequencies that allow managers to demonstrate the effectiveness and return on investment of preventive mechanisms. Without hard evidence of attack frequencies, security managers are left in the position of the man waving a dead chicken around his head while standing on a street corner. Asked why he is doing that, he answers, "To keep the flying elephants away." "But there are no flying elephants," protest the observers. "See? It works!" replies the lunatic. For more information on auditing, see Chapter 54 in this *Handbook.*

Although intrusion detection and intrusion prevention systems are necessary as system security functions, they are not sufficient to protect systems from all security threats. IDS and IPS must be a part of a more comprehensive security strategy that includes vulnerability assessment, security policy and procedural controls, network firewalls, strong identification and authentication mechanisms, access control mechanisms, file and link encryption, file integrity checking, physical security measures, and security training.

27.1.4 Brief History of Intrusion Detection. Intrusion detection is the automation of manual processes that originated in the earliest days of data processing. Joseph Wassermann of the Bell Telephone Company documented the origin of system and security audit as early as the mid-1950s, when the first computerized business system was being designed and implemented.[2]

Auditability was a key security feature from the earliest days of computer security, as proposed in J. P. Anderson's 1973 research study chartered by the U.S. Air Force.[3] Anderson proposed a scheme for automating the review of security audit trails in 1980, in a research report considered by many to be the seminal work in intrusion detection.[4] Dorothy Denning and Peter Neumann led a study of intrusion detection, conducted from 1984 to 1986, producing another seminal work in intrusion detection in 1986, in which Denning proposed a model for intrusion detection.[5]

An instantiation of Denning's intrusion detection model was prototyped as the Intrusion Detection Expert System (IDES) by a team at SRI International. IDES was a hybrid system that constructed statistical profiles of user behaviors as derived from operating system kernel audit logs and other system data sources. IDES also provided a rules-based expert system that allowed users to specify patterns of events to be flagged as intrusions.[6] IDES and the Next Generation IDES (NIDES) system that followed it marked an era in which numerous intrusion detection research projects and prototype systems were developed, including Haystack (Haystack Labs and U.S. Air Force), NADIR (Los Alamos National Laboratory), Wisdom and Sense (Los Alamos National Laboratory and Oak Ridge National Laboratory), ISOA (PRC, Inc.), TIM (Digital Equipment Corporation), ComputerWatch (AT&T), and Discovery (TRW, Inc).[7]

In the late 1980s, researchers at the University of California, Davis, designed the first network-based intrusion detection system (initially called the Network Security Monitor, but later renamed NID), which functioned much the same as many current commercial network-based intrusion detection products.[8] A subsequent U.S. Air Force–funded research product called the Distributed Intrusion Detection System (DIDS), explored coordinating network-based and host-based intrusion detection systems. DIDS was prototyped by teams at the University of California, Davis, Haystack Laboratories, and Lawrence Livermore National Laboratory.[9]

Intrusion prevention was proposed as a logical next step to intrusion detection almost from the start of intrusion detection research. Support for certain models of intrusion prevention grew when concerns regarding attacks on the TCP/IP network infrastructure (e.g., packet flooding and malformed packet attacks) grew in the mid- to late-1990s. Other types of IPS were proposed to deal with kernel-level hacks and information leakage issues.

27.2 MAIN CONCEPTS. Several strategies used in performing intrusion detection serve to describe and distinguish specific intrusion detection systems. These affect the threats addressed by each system and often prescribe the environments in which

specific systems should be used. As noted, intrusion prevention relies first on intrusion detection strategies; thus, the differentiators will be highlighted as appropriate.

27.2.1 Process Structure. Intrusion detection is defined as a monitoring and alarm generation process, and, as such, it can be described using a simple process model. This model is outlined here and will be used to illustrate the fundamental concepts of intrusion detection.

27.2.1.1 Information Sources. The first stage of the intrusion detection process comprises one or more information sources, also known as event generators. Information sources for intrusion detection may be categorized by location: network, host, or application.

27.2.1.2 Analysis Engine. Once event information is collected, it is passed to the next stage of the intrusion detection process, in which it is analyzed for symptoms of attack or other security problems.

27.2.1.3 Response. When the analysis engine diagnoses attacks or security problems, information about these results is revealed via the response stage of the intrusion detection process. Responses span a wide spectrum of possibilities, ranging from simple reports or logs to automated responses that disrupt attacks in progress. The presence of these automated responses defines an intrusion prevention system.

27.2.2 Monitoring Approach. The first major classifier used to distinguish intrusion detection systems is the monitoring approach of the system. Monitoring is the action of collecting event data from an information source and then conveying that data to the analysis engine.

The monitoring approach describes the perspective from which intrusion detection monitoring is performed. The primary monitoring approaches found in intrusion detection systems today are *network based*, *host based*, and *application based*.

27.2.3 Intrusion Detection Architecture. Even in the early days of manual security audit, researchers noted that in order for audit information to be trusted, it should be stored and processed in an environment separate from the one monitored. This requirement has evolved to include most intrusion detection approaches, for three reasons:

1. To keep an intruder from blocking or nullifying the intrusion detection system by deleting information sources
2. To keep an intruder from corrupting the operation of the intrusion detector in order to mask the presence of the intruder
3. To manage the performance and storage load that might result from running intrusion detection tasks on an operational system

In this architecture, the system running the intrusion detection system is called the *host.* The system or network being monitored is called the *target.*

27.2.4 Monitoring Frequency. Another common descriptor for intrusion detection approaches is the timing of the collection and analysis of event data. This is

usually divided between *batch-mode* (also known as *interval-based*) and *continuous* (also known as *real-time*) approaches.

In batch-mode analysis, the event data from the information source are conveyed to the analysis engine in a file or other block form. As the name suggests, the events corresponding to a particular interval of time are processed (and results provided to the user) after the intrusion has taken place. This model was the most common for early intrusion detection because system resources did not allow real-time monitoring or analysis.

In real-time analysis, event data from the information source are conveyed to the analysis engine as the information is gathered. The information is analyzed immediately, providing the user with the opportunity to respond to detected problems quickly enough to affect the outcome of the intrusion.

27.2.5 Analysis Strategy. In intrusion detection, there are two prevalent analysis strategies, *misuse detection* and *anomaly detection.*

In misuse detection, the analysis engine filters event streams, matching patterns of activity that characterize a known attack or security violation. In anomaly detection, the analysis engine uses statistical or other analytical techniques to spot patterns corresponding to abnormal system use. Anomaly detection is based on the premise that intrusions significantly differ from normal system activity. In general, intrusion detection systems rely more heavily on anomaly detection and quantitative measures to detect and block attacks.

27.3 INTRUSION PREVENTION. As discussed, IPSs are often considered a special case of IDS in which automated responses are specified. However, with the adoption of the first generation of network IPS products, additional specifications for IPSs have evolved in addition to those assigned to IDS.

27.3.1 Intrusion Prevention System Architecture. As in intrusion detection, most IPSs separate the monitoring and analysis platform from the target platform being monitored. Additional differentiations are drawn between those IPSs that separate the monitoring and analysis platform from the response platform (these are labeled "stand-alone" IPSs) and those that integrate all the functions in a single unit, usually a firewall, network switch, or router (these are labeled "integrated" IPSs.)

27.3.2 Intrusion Prevention Analysis Strategy. IPSs generally use the same structural approach to data analysis as IDS, but the nomenclature for the analysis strategies differs. IPS analysis schemes fall into two general categories, rate based and content based.

Rate-based IPS analysis makes the decision to block network traffic based on indicators of network load, as measured by statistics such as connect rates and connection counts. This category of analysis is especially useful for detecting packet flood distributed denial-of-service (DDoS) attacks.

Content-based IPS analysis makes the decision to block network traffic based on indicators of anomalous packets and specific content (often represented as IDS signatures). This approach is useful for detecting malformed packet DDoSs and other types of attacks not readily spotted by quantitative measures.[10]

27.4 INFORMATION SOURCES. Information sources represent the first stage of the intrusion detection process. They provide event information from monitored

systems on which the intrusion detection process bases its decisions. Information sources encompass both *raw* event data (e.g., data collected directly from system audit and logging mechanisms) as well as data output by system management utilities (e.g., file integrity checkers, vulnerability assessment tools, network management systems, and even other intrusion detection systems). In this section, information sources for intrusion detection are classified by location: network, host, or application.

27.4.1 Network Monitoring. The most common monitoring approach utilized in intrusion detection systems is *network based*. In this approach, information is gathered in the form of network packets, often using network interface devices set to promiscuous mode. (Such a device operating in promiscuous mode captures all network traffic accessible to it—usually on the same network segment—not just traffic addressed to it.) Other approaches for performing network-based monitoring include the use of *spanning ports* (specialized monitoring ports that allow capture of network traffic from all ports on a switch) on network switches or specialized Ethernet network taps (e.g., sniffers) to capture network traffic.

27.4.2 Operating System Monitoring. Some monitors collect data from sources internal to a computer. These differ from network-based monitoring in the level of abstraction at which the data is collected. *Host-based* monitoring collects information from the operating system (OS) level of a computer. The most common sources of operating system–level data are operating system audit trails, which are usually generated within the OS kernel, and system logs, which are generated by OS utilities.

27.4.3 Application Monitoring. *Application-based* monitoring collects information from running software applications. Information sources utilized in application-based approaches include application event logs and application configuration information.

Application-based information sources are steadily increasing in importance as systems complexity increases. The advent of object-oriented programming techniques introduces data object naming conventions that nullify much of an analyst's ability to make sense of file access logs. In this situation, the application level is the only place in the system in which one can "see" the data accesses at an appropriate level of abstraction likely to reveal security violations.

One special case of application-based monitoring comprises an entire product category in security. This type of system (sometimes called extrusion detection) monitors data transfers, looking for anomalies associated with data movement across policy boundaries. Such data monitoring is very popular as a compliance mechanism for enterprises dealing with consumer, financial, or other regulated data.

27.4.4 Other Types of Monitoring. As noted, intrusion detection information sources are not limited to raw event data. In fact, allowing intrusion detection systems to operate on results from other systems often optimizes the quality of the intrusion detection system's results. When the data are provided by other parts of the system security infrastructure (e.g., network firewalls, file integrity checkers, virus scanners, or other intrusion detection systems), the sensitivity and reliability of the intrusion detection system's results can increase significantly.

27.4.5 Issues in Information Sources. There are several issues involving information sources for intrusion detection. The major ones that have persisted over the history of intrusion detection include:

- In host-based systems, there must be a balance between collecting enough information to accurately portray security violations and collecting so much information that the collection process cripples the monitored system.
- The fidelity of the intrusion detection process is dependent not only on collecting the appropriate information but on collecting it from appropriate vantage points within the monitored system or network.
- If the IDS is expected to produce event records that will be used to support legal processes, the system must collect and handle event information in a way that complies with legal rules of evidence.
- The information collected by IDSs often includes information of a sensitive nature. This information must be secured and handled in a way that complies with legal and ethical standards.

27.5 ANALYSIS SCHEMES. Once information sources and sensors are defined and placed, the information so gathered must be analyzed for signs of attack. The *analysis engine* serves this purpose in intrusion detection, accepting event data from the information source and examining it for symptoms of security problems. As mentioned earlier, intrusion detection systems typically provide analysis features that fall into two categories, *misuse detection* and *anomaly detection.*

27.5.1 Misuse Detection. *Misuse detection* is the filtering of event streams for patterns of activity that reflect known attacks or other violations of security policy. Misuse detectors use various pattern-matching algorithms, operating on large databases of attack patterns or *signatures.* Most current commercial intrusion detection systems support misuse detection.

Misuse detection presumes that there is a clear understanding of the security policy for the system, which can be expressed in patterns corresponding to desirable activity and undesirable activity. Therefore, signatures can be described in terms of "this should never happen" as well as "only this should ever happen." Signatures also can range from simplistic *atomic* (one-part) checks to rather complex *composite* (multipart) checks. An example of an atomic check is a buffer overflow signature, in which one looks for a particular command, followed by a string exceeding a particular length. An example of a composite check is a race condition signature, in which a series of carefully timed commands occur. Signatures are gathered and structured in some way to optimize the filtering of event data against them.

The next requirement for misuse detection is that the event data collected from information sources be encoded in a way that allows it to be matched against the signature data. There are various ways of doing this, ranging from regular expression matching (sometimes called "dirty word" matching) to complex coding schemes involving state diagrams and Colored Petri Nets. State diagrams are a graphical scheme for modeling intrusions. They express intrusions in terms of *states*, represented by nodes or circles, and *transitions,* represented by lines or arcs. Colored Petri Nets are an extension of the state diagram technique that add colored *tokens,* which occupy state nodes, and whose color expresses information about the context of the state.

In some IDSs and content-based IPSs, significant resources are devoted to identifying malformed network packets, especially those in which the format of the content of the packet does not match the format of the service (e.g., SMTP) or of the associated port number of the packet. This malformed packet scheme represents one of the major categories of DDoS attacks, which seek to deny network access to legitimate users.

27.5.2 Anomaly Detection. *Anomaly detection* is the analysis of system event streams, characterizing them using statistical and other classification techniques in order to find patterns of activity that appear to deviate from normal system operation. This approach is based on the premise that attacks, and other security policy violations, are a subset of abnormal system events.

Several common techniques are used in anomaly detection:

- **Quantitative analysis.** Most modern systems that use anomaly detection provide quantitative analysis, in which rules and attributes are expressed in numeric form. The most common forms of quantitative analysis are triggers and thresholds, in which system attributes are expressed as counts occurring during some time interval, with some level defined as permissible. Triggers and thresholds can be simple, in which the permissible level is constant, or heuristic, in which the permissible level is adapted to observed levels. Network intrusion prevention systems targeting DDoS attacks often use heuristic triggers and thresholds to characterize the normal bandwidth loads, connect rates, and connection counts in network traffic.
- **Statistical analysis.** Most early anomaly detection systems used statistical techniques to identify abnormal data. In statistical analysis, profiles are built for each user and system resource, and statistics are calculated for a variety of user and resource attributes for a particular interval of time (usually a "session," defined as the time elapsed between login and logout).
- **Learning techniques.** There has been a great deal of research interest in using various learning techniques, such as neural networks and fuzzy logic, in performing anomaly detection. Despite encouraging results, there remain many practical impediments to using these techniques in production environments. The practical impediments arise due to the mismatch between those attributes that are suitable for characterization by neural networks and fuzzy logics and those attributes that are actionable by operational personnel and systems. The value of using neural networks (most of them utilizing fuzzy logic) is that they can characterize and recognize very subtle signs of trouble in systems. This is of value in situations where the problems being detected are not subtle. However, in security breaches, the difference between normal behavior and security attack is often influenced by the system context; in these scenarios, few, if any, neural networks can provide insights regarding how they reached their decisions regarding the suspicious events on which they trigger. A security person in an operational context usually requires this sort of insight in order to devise an appropriate reaction to the detected intrusion.
- **Advanced techniques.** Anomaly detection as applied to intrusion detection remains an active research area. Recent research efforts include the application of such advanced analytic techniques as genetic algorithms, data mining, autonomous agents, and immune system approaches to the problem of recognizing new attacks

and security violations. Again, these techniques have not yet been widely fielded in commercial IDSs, although they have appeared in special purpose products.

27.5.3 Hybrid Approaches. There are significant issues associated with both misuse detection and anomaly detection approaches to event analysis for intrusion detection; however, combining both approaches provides considerable benefit. The anomaly detection engine can allow the IDS to detect new or unknown attacks or policy violations. This is especially valuable when the target system protected by the IDS is highly visible on the Internet or other high-risk network. In IPS, anomaly detection applied to network traffic attributes allows one of the only means to deal with packet-flood DDoS attacks, a growing concern in today's networks.

The misuse detection engine, in turn, protects the integrity of some anomaly detection engines by assuring that a patient adversary cannot gradually change behavior patterns over time in order to retrain the anomaly detector to accept attack behavior as normal. Thus, the misuse detector mitigates a significant deficiency of anomaly detection for security purposes.

27.5.4 Issues in Analysis. Here are a few of the many issues in intrusion detection analysis:

- Misuse detection systems, although very effective at detecting those scenarios for which detection signatures have been defined, cannot detect new attacks.
- Anomaly detection systems are capable of detecting new attacks but usually have false positive rates so high that users often ignore the alarms they generate.
- Anomaly detection systems that rely on artificial intelligence (AI) techniques often suffer from a lack of adequate training data. (Data are used to define the detector's logic for distinguishing "normal" from "abnormal" events.)
- Malefactors with access privileges to the system, while anomaly-detection systems are being trained, can covertly teach the system to accept specific patterns of unauthorized activities as normal. Later, the anomaly-detection systems will ignore the actual misuse.

27.6 RESPONSE. The final stage of intrusion detection, response, consists of the actions taken in response to the security violations detected by the IDS. Responses are divided into *passive* and *active* options. The difference between passive and active responses is whether the user of the IDS, or the system itself, is responsible for reacting to the detected violations. As discussed, the former option is associated with classic IDS; the latter is associated with IPS.

27.6.1 Passive Responses. When passive responses are selected, the IDS simply provides the results of the detection process to the user, who must then act on these results, independent of the IDS. In this option, the user has total control over the response to the detected problem. In some IDSs, the information provided to the user regarding detection results can be divided into *alarms* and *reports.*

27.6.1.1 Alarms. Alarms are messages that are communicated immediately to users. Commercial IDSs use a variety of channels for conveying these alarms to security personnel. The most common is a message screen or icon written to the IDS control

console. Other alarm channels include pagers, email, wireless messaging, and network management system traps.

27.6.1.2 Reports. Reports are messages or groups of messages that are generated on a periodic basis. They typically document events that have happened in the past and often include aggregate figures and trends information. Many commercial IDS products support extensive reporting features, allowing a user to set up automatic report generation with several versions, each targeting a different level of management.

27.6.2 Active Responses: Man-in-the-Loop and Autonomous. When an IDS provides active response options, these usually fall into two categories. The first requires the IDS to take action, but with the active involvement of an interactive user. This option is sometimes called a man-in-the-loop mechanism. This option is preferred for critical systems, as it allows an operator to track an attacker or intervene in a sensitive situation in a flexible, exacting way.

The other active response option, which usually defines an IPS, provides for preprogrammed actions taken automatically by the system with no human involvement. The automated response option is required when dealing with certain sorts of automated attack tools (viruses or worms) or DDoS attacks. These attacks proceed at machine speed and therefore are outside the reach of a human-controlled manual intervention. As automated and DDoS attacks are the tool of choice for online extortionists, the number of IPSs fielded in commercial networks has grown rapidly over the past few years.

27.6.3 Automated Response Goals. Automated responses support three categories of response goals:

1. Collecting more information about the intrusion or intruder
2. Amending the environment (e.g., changing a switch or router setting to deny access to an intruder)
3. Taking action against the intruder

Although the last of these groups, sometimes labeled *strike back* or *hack back*, occasionally is discussed in security circles, the other options are far more productive in most situations. At this time, taking action against the intruder is considered inappropriate in almost all situations, and should be undertaken only with the advice and counsel of a legal authority.

Amending the environment and collecting more information can occur in either stand-alone or integrated fashions.

27.6.3.1 Stand-Alone Responses. Some automated responses are designed to use features that fall entirely within the intrusion detection system. For instance, an intrusion detection system may have special detection rules, more sensitive or detailed than those provided in normal modes of operation. In a stand-alone adaptive response, the IDS would use the more sensitive rules when evidence of the preamble of an attack is detected. This allows the IDS to turn sensitivity levels up only when the additional detection capabilities are needed, so that false alarm rates are reduced.

27.6.3.2 Integrated Responses. The response option often considered the most productive is that of using integrated measures that change the system settings to block the attacker's actions. Such responses can affect the configuration of the target system, the IDS/IPS host, or the network on which both reside. In the first case, the IDS/IPS might change the settings of the logging mechanisms on the target host to increase the amount or type of information collected. The IDS also might change its analysis engine so that more subtle signs of attack are recognized. In another response option reflected in commercial products, the IDS responds to an observed attack signature by querying the target system to determine whether it is vulnerable to that specific attack. Should the vulnerability be present, the IDS directs the target system to correct that vulnerability. In effect, this process provides the target system with an immune function and permits it to "heal" itself, either blocking an attack outright or else interactively repairing any damage done in the course of the attack. Finally, some systems may use special-purpose decoy systems, called *honey pots* or *padded cells,* as diversions for attackers. When these systems are provided, the IDS may be configured to divert attackers into the decoy environments.

In a special case of integrated response seen in many commercial IPS offerings, the IPS is integrated with a switch or router. When an attack is detected, the switch or router is reconfigured on the fly to block the source of the attack. Other IPS offerings that are designed to deal with DDoS attacks use multiple IDS/IPSs to detect the attacks, then manipulate the switching fabric[11] to divert the attacks from the targeted systems. Over time, such IPS features will likely be integrated with network infrastructure devices, as many firewall features already have done.

27.6.4 Investigative Support. Although the primary design objective of intrusion detection systems is detecting attacks and other possibly problematic system events, information collected and archived by IDSs also can support those charged with investigating security incidents. This functional requirement may levy additional technical requirements on IDSs. For instance, if investigators plan to use IDS monitoring features to perform a targeted surveillance of an attack in progress, it is critical that the information sources be "silent" so that adversaries are not aware that they are being monitored. Furthermore, the IDS monitors must be able to convey information to the investigators through a trustworthy, secure channel. Finally, the IDS itself must be under the control of the investigators or other trusted parties; otherwise, the adversaries may mask their activities by selectively spoofing information sources. Perhaps the most important thing for investigators to remember about IDSs is that the information provided should be corroborated by other information sources (e.g., network infrastructure device logs), not necessarily accepted at face value.

27.6.5 Issues in Responses. As in information sources and analysis strategies, certain issues associated with IDS response features have endured over the history of intrusion detection. The principal issues are:

- Users' needs for IDS response capabilities are as varied as the users themselves. In some systems environments, the IDS response messages are monitored around the clock, with real-time action taken by system administrators based on IDS alarms. In other environments, users may use IDS responses, in the form of reports, as a metric to indicate the threat environment in which a particular system resides. It is important to consider the specific needs of the user when selecting an IDS.

- Given false-positive error rates for IDSs, response options must be tunable by users. Otherwise, users will simply tune out the IDS responses. This nullifies the value of the IDS.
- When the IDS provides automated responses to detected problems, there is a risk of the IDS itself launching an effective denial-of-service attack against the system it is protecting. For instance, suppose an IDS is configured with rules that tell it "upon detecting an attack from a given IP address, direct the firewall to block subsequent access from that IP address." An attacker, knowing this IDS is so configured, can launch an attack with a forged IP source address that appears to come from a major customer or partner of the organization. The IDS will recognize the attack and then block access from that organization for some period of time, effecting a denial of service.

27.7 NEEDS ASSESSMENT AND PRODUCT SELECTION. The value of intrusion detection products within an organization's security strategy is optimized by a thorough needs assessment. These needs and security goals can be used to guide the selection of products that will enhance the security stance of the organization.

27.7.1 Matching Needs to Features. The needs most often addressed by intrusion detection and intrusion prevention systems include:

- Prevention of problem behaviors by increasing the risk of discovery and punishment for system attackers.
- Detection of security violations not prevented (or even, in some cases, not preventable) by other security measures.
- Documentation of the existing level of threat to an organization's computer systems and networks.
- Detection and, where possible, mitigation of attack preambles. (These include activities such as network probes, port scans, and other such "doorknob rattling.")
- Diagnosis of problems in other elements of the security infrastructure (i.e., malfunctions or faulty configurations).
- Granting system security personnel the ability to test the security effects of maintenance and upgrade activities on the organizational networks.
- Providing information about those violations that do take place, enabling investigators to determine and correct the root causes.
- Providing evidence of compliance with a given regulatory requirement for information protection. This represents a significant need for members of various regulated industries, such as banking and health care.

Regardless of which of these specific needs are relevant to the user, it is important to consider the ability of the intrusion detection system to satisfy the needs of the specific environment in which it is installed. A critical part of this determination is considering whether the intrusion detection system has the ability to monitor the specific information sources available in the target environment. What is even more important is whether the organizational security policy translates into a monitoring and detection policy that can be used to configure the IDS (or in the case of an IPS, a monitoring, detection, and response policy.) The structure of the security policy is especially critical to the success of an IPS.

27.7.2 Specific Scenarios. There is no universally applicable description for computer networks or the IDSs that protect them. There are, however, some common scenarios, given current trends in networking and system usage.

A popular justification for using IDSs early in an organization's security life cycle is to establish the threat level for a given network enclave. Network-based IDSs often are used for this purpose, with monitors placed outside the organizational firewall. Those who are responsible for winning management support for security efforts often find this use of IDSs to be quite helpful.

Many organizations use IDSs to protect Web servers. In this case, the nature of the interactions that the Web server has with users will affect the selection and configuration of the IDS. Most Web servers serve two types of functions: (1) informational (e.g., Web servers that support simple HTTP and FTP queries from users) and (2) transactional (e.g., Web servers that allow user interaction beyond simple HTTP or FTP traffic). Transactional Web servers are usually more difficult to monitor than informational servers, as the range of interactions between users and servers is wider. For critical transactional Web servers, security managers may wish to consider multiple IDSs, monitoring the servers at multiple levels of abstraction (i.e., application, host, and network).

The third scenario involves organizations that wish to use IDSs as additional protection for specific portions of their networked systems. An example of this is the medical organization that wishes to protect the patient record database systems from privacy breaches. In this situation, as in the Web server example just given, it may be advisable to use multiple IDSs, monitoring interactions at multiple levels of abstraction. The output of these multiple systems can be synchronized and inconsistencies noted for a reliable indication of threat levels. Another example that is increasingly common is the organization that is concerned about wireless connectivity. In this case, WiFi monitoring products are commercially available, with information collection and monitoring features that are similar to those of classic IDSs.

In recent years, the expanding use of wireless local area networks (WLANs) in buildings and campuses has stimulated the development of wireless intrusion detection and prevention systems (WIDPSs). The basic principles are the same as for other IDSs and IPSs, with the addition of an interesting wrinkle: The WIDPSs are often used to discover unauthorized access points installed on an organization's WLANs by rogue employees or by intruders. Karen Scarfone and Peter Mell, in their July 2012 edition of NIST SP 800-94, point out that WIDPS sensors should be placed in areas where there should be no wireless network activity. Some security managers walk through and drive around their facilities with WIDPS tools on their laptop computers to identify unauthorized access points. However, completely passive eavesdropping on WLAN traffic cannot be detected.[12]

In extensive networks, should the security architect decide to layer multiple IDSs and IPSs, a security event monitor/security information manager (SIM/SEM) may be required. Such a system would be necessary to consolidate and integrate the results of each IDS/IPS into a coherent set of conclusions.

27.7.3 Integrating IDS Products with Your Security Infrastructure. As mentioned, an IDS is not a substitute for a firewall, virtual private network, identification and authentication package, or any other security point product. However, an IDS can improve the quality of protection afforded by the other point products by monitoring their operation, noting signs of malfunction or circumvention. Furthermore,

an IPS can interact in concert with the rest of the point products to help block an attack in progress.

27.7.4 Deployment of IDS Products. The first generations of IDS installations have yielded some insights associated with deployment of IDSs. The key points include the location of sensors, scheduling the integration of IDSs, adjusting alarm settings, and outsourcing IDS/IPS services.

27.7.4.1 Location of Sensors. There are four general locations for IDS sensors:

1. Outside the main organizational firewall
2. In the network DMZ (inside the main firewall, but outside the internal firewalls)
3. Behind internal firewalls
4. In critical subnets, where critical systems and data reside

As mentioned, IDS sensors placed outside the main organizational firewall are useful for establishing the level of threat for a given network. Sensors placed within the DMZ[13] can monitor for penetration attempts targeting Web servers. IDSs monitors for internal attacks are placed on internal network segments, behind internal firewalls. And for critical subnets, IDS sensors usually are placed at the choke points at which the subnets are connected to the rest of the corporate network. In the case of wireless networking, specialized IPS devices serve to discover and report unauthorized access to networks via WLAN access points or open software access points, through misconfigured WLAN interfaces on laptops and other WLAN devices. These IPS devices are usually placed behind firewalls and distributed across the physical space occupied by the organization and its users.

27.7.4.2 IDS Integration Scheduling. Early generations of intrusion detection products proved that integration processes must not be rushed. IDSs still rely on operator interactions to screen out false alarms and to act on legitimate alarms. Hence, it is critical that the processes provide adequate time for operational personnel to learn the behavior of the IDS on target systems, developing a sense of how the IDS interoperates with particular system components in different situations. This wisdom applies even more to IPS installation, where a miscue in specifying a response can have disastrous effects on the function of critical networks.

27.7.4.3 Alarm Settings. IDSs have significant false alarm rates, with false positive rates as high as 80 percent in some situations. Many knowledgeable IDS integrators advise that alarms be suspended for a period of weeks, even months, as operators gain familiarity with the IDS and target systems. It is especially wise to delay activation of automated responses to attacks until operators and system administrators are familiar with the IDS and have tuned it to the target environment.

27.7.4.4 Outsourcing of IDS/IPS. Any discussion of IDS and IPS strategies would be incomplete without mention of outsourcing these security services. There are significant advantages associated with this approach, especially in enterprises too small to afford an extensive security staff. As in other areas of IT outsourcing, one must have an extremely clear idea of the specific security and operational goals desired

in order for this approach to be effective. A clearly worded security policy that reflects current concerns is essential. The advantages of outsourcing are many: Managed security service providers usually have considerable experience in dealing with IDS and IPS equipment, their staffs are often well trained and experienced in the use of the equipment, monitoring personnel are usually in attendance around the clock, and contract terms often include specific levels of service agreements.

In the words of a wise CISO, "outsourcing isn't offloading." That means that outsourcing security functions does not relieve you of the responsibility for system security. It also means that you must exercise due diligence in selecting the service provider, tasking and managing it, and monitoring it to ensure that your policy goals are being well served by the provider.

27.8 CONCLUSION. Intrusion detection and intrusion prevention are valuable additions to system security suites, allowing security managers to spot, and sometimes block, those security violations that inevitably occur despite the placement of preventive security measures. Although current commercial products are imperfect, they serve to recognize many common intrusion types, in many cases quickly enough to allow security personnel and IPSs to block damage to systems and data. Furthermore, as research and development in intrusion detection and prevention continues, the quality and capabilities of available IDSs and IPSs will steadily improve.

27.9 FURTHER READING

Crosbie, M., and E. H. Spafford. "Defending a Computer System Using Autonomous Agents." Proceedings of the 18th National Information Systems Security Conference. Baltimore, MD: October 1995.

Jackson, K. A., D. DuBois, and C. Stallings. "An Expert System Application for Network Intrusion Detection." Proceedings of the 14th National Computer Security Conference. Washington, DC: October 1991.

Fadia, A., and M. Zacharia. *Intrusion Alert: An Ethical Hacking Guide to Intrusion Detection*. Vikas Publishing House, 2010.

Flegel, U. *Privacy-Respecting Intrusion Detection*. New York: Springer, 2007.

Hämmerli, B., and R. Sommer, eds. *Detection of Intrusions and Malware, and Vulnerability Assessment: 4th International Conference*, DIMVA 2007 Lucerne, Switzerland, July 12–13, 2007 Proceedings. New York: Springer, 2007.

Kumar, Sandeep, and E. Spafford. "A Pattern Matching Model for Misuse Intrusion Detection." Proceedings of the 17th National Computer Security Conference. Baltimore, MD: October 1994.

Lunt, T., *et al.* "A Real-Time Intrusion Detection Expert System (IDES)." *Computer Science Lab*, Menlo Park, CA: SRI International, May 1990.

Mukherjee, B., L. T. Heberlein, and K. N. Levitt. "Network Intrusion Detection," *IEEE Network* 8, No. 3 (May–June 1994).

Paxson, V. "Bro: A System for Detecting Network Intruders in Real Time." *7th USENIX Security Symposium*. San Antonio, TX: January 1998.

Porras, P., and P. Neumann. "EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances." *Proceedings of 20th National Information System Security Conference*. Baltimore, MD: October 1997.

Scarfone, K., and P. Mell. *Guide to Intrusion Detection and Prevention Systems (IDPS)* Revision 1 (Draft). NIST Special Publication 800-94.

Gaithersburg, MD: National Institute of Standards and Technology, 2012. Available: http://csrc.nist.gov/publications/drafts/800-94-rev1/draft_sp800-94-rev1.pdf

Schaefer, M., et al. "Auditing: A Relevant Contribution to Trusted Database Management Systems." *Proceedings of the 5th Annual Computer Security Applications Conference*. Tucson, AZ: December 1989.

Shostack, A., and S. Blake. "Towards a Taxonomy of Network Security Assessment Techniques." *Proceedings of 1999 Black Hat Briefings*. Las Vegas, NV: July 1999.

Trost, R. *Practical Intrusion Analysis: Prevention and Detection for the Twenty-First Century*. Addison-Wesley Professional, 2009

27.10 NOTES

1. See the Telecom Glossary 2000 from the American National Standards Institute, Inc.: www.its.bldrdoc.gov/projects/telecomglossary2000 (URL inactive).
2. J. J. Wassermann, "The Vanishing Trail," *Bell Telephone Magazine* 47, No. 4 (July/August 1968).
3. J. P. Anderson, "Computer Security Technology Planning Study Volume II," ESD-TR-73-51, Electronic Systems Division, Air Force Systems Command, Hanscom Field, Bedford, MA 01730 (October 1972).
4. J. P. Anderson, *Computer Security Threat Monitoring and Surveillance* (Fort Washington, PA: James P. Anderson Co., April 1980).
5. D. Denning, "An Intrusion Detection Model," *Proceedings of the 1986 IEEE Symposium on Security and Privacy* (Washington, DC: IEEE Computer Society Press, 1986).
6. T. Lunt et al., "A Real-Time Intrusion Detection Expert System (IDES)," *Computer Science Lab*, SRI International, May 1990.
7. B. Mukherjee, L. T. Heberlein, and K. N. Levitt, "Network Intrusion Detection," *IEEE Network* 8, No. 3 (May–June 1994).
8. L. T. Heberlein, K. N. Levitt, and B. Mukherjee, "A Network Security Monitor," *Proceedings of the 1990 IEEE Symposium on Research in Security and Privacy,* Oakland, CA: May 1990.
9. S. Snapp et al., "DIDS (Distributed Intrusion Detection System) Motivation, Architecture, and an Early Prototype," *Proceedings of the 14th National Computer Security Conference,* Washington, DC: October 1991.
10. IPS Quadrant, *Information Security Magazine* (July 2004); http://infosecuritymag.techtarget.com/ss/0,295796,sid6_iss426_art880,00.html
11. A. Freedman, *Computer Desktop Encyclopedia*, v21.1, 2008 (www.computerlanguage.com) provides this definition: "Switch fabric—(1) The interconnect architecture used by a switching device, which redirects the data coming in on one of its ports out to another of its ports. The word 'fabric' comes from the resulting criss-crossed lines when all the inputs on a switch with hundreds of ports are connected to all possible outputs. (2) The combination of interconnected switches used throughout a campus or large geographic area, which collectively provide a routing infrastructure."
12. K. Scarfone and P. Mell, *Guide to Intrusion Detection and Prevention Systems (IDPS)*, Revision 1 (Draft). NIST Special Publication 800-94

(Gaithersburg, MD: National Institute of Standards and Technology, July 2012); http://csrc.nist.gov/publications/drafts/800-94-rev1/draft_sp800-94-rev1.pdf, pp. 36–48.

13. The DMZ is a reserved area in some network architectures in which Web servers are often placed, separated from the Internet by one firewall system and separated from the internal corporate network by another firewall.

CHAPTER 28

IDENTIFICATION AND AUTHENTICATION

Ravi Sandhu, Jennifer Hadley, Steven Lovaas, and Nicholas Takacs

28.1 INTRODUCTION 28·2

28.2 FOUR PRINCIPLES OF AUTHENTICATION 28·2

28.2.1 What Only You Know 28·3
28.2.2 What Only You Have 28·3
28.2.3 What Only You Are 28·4
28.2.4 What Only You Do 28·5

28.3 PASSWORD-BASED AUTHENTICATION 28·5

28.3.1 Access to User Passwords by System Administrators 28·5
28.3.2 Risk of Undetected Theft 28·6
28.3.3 Risk of Undetected Sharing 28·7
28.3.4 Risk of Weakest Link 28·7
28.3.5 Risk of Online Guessing 28·8
28.3.6 Risk of Off-Line Dictionary Attacks 28·9
28.3.7 Risk of Password Replay 28·10
28.3.8 Risk of Server Spoofing 28·11
28.3.9 Risk of Password Reuse 28·12
28.3.10 Authentication Using Recognition of Symbols 28·13

28.4 TOKEN-BASED AUTHENTICATION 28·13

28.4.1 Card Entry Systems 28·13
28.4.2 Proximity and Touch Cards 28·14
28.4.3 Smart Cards and Dongles 28·15
28.4.4 Soft Tokens 28·16
28.4.5 One-Time Password Generators 28·17
28.4.6 Authentication Using Mobile Devices 28·17

28.5 BIOMETRIC AUTHENTICATION 28·18

28.6 CROSS-DOMAIN AUTHENTICATION 28·18

28.7 RELATIVE COSTS OF AUTHENTICATION TECHNOLOGIES 28·18

28.8 CONCLUSIONS 28·18

28.9 SUMMARY 28·19

28.10 FURTHER READING 28·20

28.11 NOTES 28·21

28.1 INTRODUCTION. *Authorization* is the allocation of permissions for specific types of access to restricted information. In the real world, authorization is conferred on real human beings; in contrast, information technology normally confers authorization on *user identifiers* (IDs). Computer systems need to link specific IDs to particular authorized users of those IDs. Even inanimate components, such as network interface cards, firewalls, and printers, need IDs. *Identification* is the process of ascribing an ID to a human being or to another computer or network component. *Authentication* is the process of *binding* an ID to a specific entity. For example, authentication of a user's identity generally involves narrowing the range of possible entities claiming to have authorized use of a specific ID down to a single person.

The focus of this chapter is on person-to-computer authentication. In practice, we also need computer-to-person authentication to prevent spoofing of services on a network. This type of authentication is increasingly important, especially on open networks such as the Internet, where users may be misled about the identity of the Websites they visit. For example, some criminals send unsolicited email messages in Hypertext Markup Language (HTML) to victims; the messages include links that are labeled to suggest an inoffensive or well-respected Website, but the underlying HTML actually links to a fraudulent site designed to trick people into revealing personal information, such as credit card numbers or details to support theft of identity. More generally, computer-to-computer mutual authentication, typically in both directions, is essential to safeguard critical transactions such as those of interbank transfers and business-to-business electronic commerce.

In the early decades of computer usage, most computers authenticated users who accessed mainframes from within a single enterprise. User IDs therefore could be assigned in a centralized and controlled manner. Even so, identifiers have never necessarily been unique, for there is no obligatory one-to-one relationship between a user ID and a human being's real-world identity. For example, several people could share an account such as *inventory clerk* without interference from the computer; at most, the operating system might be configured to prevent simultaneous sharing of an ID by limiting the number of sessions initiated with a specific ID to one.

Conversely, a single user often has many user IDs. For example, there may be unique identifiers for each of dozens of Websites for music clubs, book clubs, enterprise email, and so on. Even on the same computer, a given user might have several accounts defined for different purposes; jane_doe and jdoe might be identifiers for two different application packages on a system. These multiple identifiers cause problems for administrators if they do not know that the same user is associated with the different IDs; they also cause practical problems for users who have to use different authentication methods for a range of IDs. One of the critical goals of today's identification and authentication (I&A) research and development is to develop reliable and economical methods for *single signon*, whereby users would not have to reidentify and reauthenticate themselves when accessing different computer systems linked into an Internet. For details of I&A in facilities security, see Chapter 23 in this *Handbook*.

28.2 FOUR PRINCIPLES OF AUTHENTICATION. Authentication of a claimed identity can be established in four ways:

1. What only you know (passwords and passphrases)
2. What only you have (tokens: physical keys, smart cards)

3. What only you are (static biometrics: fingerprint, face, retina, and iris recognition)
4. What only you do (dynamic biometrics: voice, handwriting, and typing recognition)

In each approach, the assumption is that no one else but the authorized user of an identifier has access to the password or token and that the probability of simulating static or biometric data is acceptably low.

These methods can be combined; for example, passwords often are combined with tokens or biometrics to provide stronger authentication than is possible with either one alone. A familiar example of this *two-factor authentication* occurs with automatic teller machine (ATM) cards. Possession of the card (the token) and knowledge of the personal identification number (the PIN, corresponding to a password) are required to access a user's bank account.

This chapter introduces each of these four authentication methods and provides additional details for each. For discussions of methods of bypassing and subverting identification and authentication techniques, see Chapter 15 in this *Handbook*.

28.2.1 What Only You Know. Password- or passphrase-based authentication is so widely used that any person who has had any contact with computers and networks probably has had several passwords. Although password technology often is poorly administered and insecure (and frustrating) for users and administrators, passwords can be deployed much more securely and conveniently than they usually are. Many security professionals have felt and hoped for years that passwords would eventually be phased out, to be replaced by tokens or biometrics, but the consensus today is that passwords are not likely to disappear soon and that they will continue to be the dominant authentication technique for years to come.

Demonstrating knowledge of a password does not directly authenticate a human being. It simply authenticates knowledge of the password. Unauthorized knowledge of, or guessing at, a password can lead to impersonation of one user by another; this is called *spoofing*. The theft of a password can be difficult to detect since it is not a tangible asset. Passwords are also very easy to share. It is common for senior executives to give their passwords to their assistants to facilitate their work, even though assigning proxy privileges would be as effective and more secure. Students of security—including criminal hackers—quickly learn that poor password management can be widespread unless organizations provide consistent awareness, education, and training; for example, naïve users sometimes place their passwords on sticky notes under their keyboards or even in plain view on the side of their monitors.

28.2.2 What Only You Have. Authentication based on possession of a token is used where higher assurance of identity is desired than is possible by passwords alone. As with passwords, possession of a token does not directly authenticate a human being; rather it authenticates possession of the token and ability to use it. Sometimes a password or PIN is required to use the token, thus establishing two-factor authentication; the theory is that the requirement to have both elements decreases the likelihood of successful spoofing.

Tokens can take on a variety of forms. The oldest token is the physical key for a physical lock, but these are not often used for securing computer systems. *Soft tokens* are carried on transportable media or even accessed over a network from a server. Soft tokens contain only data; they typically require a password to access the contents.

Modern tokens are usually implemented in self-contained hardware with computing capability. Examples include:

- Credit card–size devices with a liquid crystal display (LCD) that display pseudorandom numbers or other codes.
- LCD devices in the shape of a key fob using the same algorithms as the credit card–shape devices.
- Hardware devices called *dongles* that plug into input-output ports on computers. Examples include dongles for serial ports, parallel ports, Universal Serial Bus (USB) ports, and PC-card interfaces.

All tokens used for computer authentication require software to process information residing in or produced by them. The most significant distinction is whether the tokens require electronic contact with the authentication system. *Contactless* tokens are easier to deploy because they do not require specialized readers. For example, the credit card and key fob pseudorandom number generators simply require the user to enter the visible code in response to a prompt from the authentication software. Contactless tokens are more limited in function than *contact tokens*. For instance, a contact token can be used to create digital signatures whereas a contactless token cannot do so practically.

In cyberspace, a token does not authenticate by means of physical characteristics. Rather the token has some secret, either exclusive to itself or possibly shared with a server on the network. Authentication of the token is really authentication of knowledge of the secret stored on the token. As such, authentication based on possession of a token is tantamount to authentication based on what the token knows. However, this secret can be longer and more random than a secret that a user has to retain in human memory, such as a password. Unfortunately, building cost-effective and secure tokens from which the secret cannot be extracted by tampering or by brute-force guesswork has proven much more difficult than initially anticipated. In the early 1990s, many security professionals believed that tokens would replace passwords; in fact, however, although tokens continue to be an attractive authentication technology, they probably will not become pervasive soon because of the consistent (but highly debatable) belief that passwords are less expensive to implement and manage than other methods of authentication (see Section 28.7).

28.2.3 What Only You Are. Biometrics takes authentication directly to the human being. This topic is covered more extensively in Chapter 29 in this *Handbook*, but the basics can be mentioned here.

As humans, we recognize each other by a number of characteristics. Biometric authentication seeks to achieve a similar result in cyberspace. A *static biometric* is a characteristic of a person such as fingerprint, hand geometry, or iris pattern; more dramatically, it could be the DNA of an individual. The likelihood of two individuals having identical fingerprints, iris patterns, or DNA is minuscule (with exceptions for genetically identical siblings). Biometrics requires specialized and expensive readers to capture the biometric data, making widespread deployment difficult.

Biometrics also suffers from the problems of replay and tampering. Thus, the biometric reader must itself be trusted and tamper-proof; this reduces the likelihood of an attacker capturing the data input and replaying it at a later time, or creating false biometric profiles to trick the system into accepting an imposter. Moreover, the biometric

data themselves must be captured in proximity to the user to reduce the likelihood of substitution, such as in the case of stolen blood used to fool a DNA-based biometric system. If the data are transmitted to a distant server for authentication, the transmission requires a secure protocol, with extensive provisions for time-stamping and rapid expiration of the data.

28.2.4 What Only You Do. *Dynamic biometrics* captures a dynamic process rather than a static characteristic of a person. A well-known example is that of signature dynamics. Signature dynamics involves recording the speed and acceleration of a person's hand as a signature is written on a special tablet. Rather than merely the shape of the signature, it is the dynamic characteristics of motion while writing the signature that authenticates the person—motions that are extremely hard to simulate. Another possibility is to recognize characteristics of a person's voice as he or she is asked to read aloud some specified text. Keystroke dynamics of a person's typing behavior is another alternative.

As in all other forms of authentication, dynamic biometrics depends on exclusion of capture and playback attacks, in which, for example, a recording of someone's voice might be used to fool a voice-recognition system. Similarly, a signature-dynamics system might be fooled by playback of the data recorded from an authentic signature. Encryption techniques help to make such attacks more difficult.

Security experts agree that biometrics offer a stronger guarantee of authentication, but deployment on a large scale remains to be demonstrated. Whether this technology becomes pervasive ultimately may be determined by its social and political acceptability as much as by improved technology.

28.3 PASSWORD-BASED AUTHENTICATION. Passwords are the pervasive technology for authentication in cyberspace today. At a conservative estimate, there are close to a billion password-based authentications per day. Examples include the vast number of Internet users and the number of passwords each one uses every day. However, the current deployment of password technology needs to be improved in many ways. Today users must remember too many identities and corresponding passwords. Also, the deployed technology is more fragile than it needs to be; for example, many users choose passwords that can be guessed easily. Passwords are never going to be as secure as the strongest biometric systems, so one would not use them as the sole basis for, say, launching nuclear missiles. However, their use can be made strong enough for many less critical transactions.

The next sections review the major risks of password use and their mitigation by technical, social, and procedural means.

28.3.1 Access to User Passwords by System Administrators. One of the most dangerous practices in use today is the storage of unencrypted user passwords accessible to system administrators. In some sites, new users receive passwords that are assigned and written down by system administrators. If these passwords are used only once, for the initial logon, the user can be forced to choose or create a truly secret password that no one else knows. However, in many such sites, administrators keep control of a paper or electronic record, usually for quick access when users forget their own passwords. Such access completely destroys an important element of I&A: *nonrepudiation.* If someone else has access to a password, then authorized users can reasonably *repudiate* transactions by claiming that their identities were spoofed. It is difficult to counter such repudiation, especially in a court of law considering

an accusation of malfeasance by the authorized user of that password. In general, passwords that will be used repeatedly should not be written down, and they should not be accessible to system administrators. Critical passwords can be written down, stored in tamper-proof containers, and locked away where at least two signatures will be required for retrieval in case of emergency.

28.3.2 Risk of Undetected Theft. Perhaps the biggest intrinsic risk with passwords is that they can be stolen without knowledge of the user. Observation of someone typing in a password is sufficient to leak it. This can happen surreptitiously without the victim's explicit knowledge. A related risk is disclosure of a password to an attacker who persuades the legitimate user to reveal it by posing as a systems administrator who needs the password to do something beneficial for the user. Loss of a physical token eventually may be discovered, since it is missing, although the possibility of cloning these devices remains. Loss of a password, however, can be discovered only by detecting its misuse or by finding it in the possession of an unauthorized user (e.g., in a list of passwords cracked by using a dictionary-based *password-cracking* program, as described in Section 28.3.6).

There are several mitigations of this risk. First, user education and awareness are critically important. People need to treat important secrets with the care they deserve. In an unsafe environment, a password should be typed in discreetly. Efforts to be discreet should be positively reinforced while negligence in exposing passwords during entry should be considered akin to bad social behavior.

User education and awareness, although extremely important, can never be the whole solution. People will inevitably slip up and make mistakes. Some of us are more negligent than others. Others will be observed surreptitiously. In some cases passwords will be revealed to computers with Trojan horses (see Chapter 16 in this *Handbook*) that capture them. Technologists must pursue technical and human solutions to mitigate these risks.

Since some losses of control over passwords are inevitable, it logically follows that password-based authentication should be used only in situations where misuse detection is not only feasible but actually convenient to do in real time. To make this possible, the system architecture should centralize the information needed for misuse detection in one place. If the required information is dispersed across many servers, it will be difficult to coordinate the different audit trails. Traditionally, users of password systems have not considered the need for misuse detection. However, modern security is firmly based on a mix of prevention and detection techniques. Security professionals should apply similar thinking to authentication systems. Ease of misuse detection should be an important criterion in the design of any authentication system. For password-based systems, misuse detection capability should be considered an essential requirement. (For information on intrusion-detection systems, see Chapter 27 of this *Handbook*.)

What else can system designers do to mitigate this risk? It should be made easy for users to change their passwords themselves. Having a system administrator change a password that will be used more than once is illogical.

If a user feels that a password may have been compromised, changing it should be a simple matter. In particular, the system should never prevent a user's attempt to change a password. Some deployed systems will deny change of a password if the password was changed recently, say in the past 24 hours. Although there are reasons for this kind of restriction, it may create a bigger risk than the one it purports to prevent.

Users should be encouraged to change their passwords fairly often; a typical allowable lifetime for a password is between 30 and 90 days. Without occasional changes,

a compromised password could be held until the malicious attacker finds opportunity to use it. Frequent changes to passwords reduce the window of opportunity for such attackers.

28.3.3 Risk of Undetected Sharing. Another major risk of passwords is the ease with which they can be shared. There are many examples of sharing between executives and their secretaries, between physicians and office staff or nurses, between professors and their secretaries or students, and among coworkers in any activity. User education and strict policies against password and account sharing are obvious first steps to deter this possibility. Strict policies can be effective within an organization, but their deterrent effect may not carry over to large consumer populations. Misuse detection also can be employed to enforce a strict policy.

The root cause of password sharing within an organization is the lack of effective delegation mechanisms whereby selected privileges of one user can be delegated to another. Better authorization mechanisms could eliminate much of the perceived need for password sharing. It should be possible for secretaries to read their bosses' email under their own identities and passwords. In fact, the bosses should be able to segregate the email that the secretaries can read while denying access to more sensitive email. Moreover, reading the boss's email should be possible without allowing the secretary to send email under the boss's identity; a proxy privilege could allow secretaries to answer their boss's email while signing the replies with their own names. In the nonelectronic world, secretaries routinely answer mail for other people without impersonating them, and this should be the practice with computers as well.

Sharing of passwords among consumers is likely to occur when the cost to consumers is minimal. Although consumers are unlikely to share passwords for an online bank or brokerage account with others, they may be willing to share passwords for an online subscription service, possibly with many friends. A dishonest consumer may even make a business of reselling the service. One way to deter such piracy would be to tie exposure of the password to exposure of a sensitive secret of the consumer, such as a credit card number. Few people, criminal or not, would hand over a password that includes their own credit card number.

Another approach that reduces account sharing is one-time passwords that are generated by inexpensive tokens that have recently been distributed to consumers by Web merchants and banks (e.g., Citibank starting in May 2006) as a method for authenticating identity.[1] These tokens generate random passwords that change every minute or so and that can be traced to the specific unit that creates them—and that unit can be tied precisely to the original recipient. Not only does such a system make password sharing virtually impossible, but a shared password can be traced directly to the violator of the terms of use and result in legal action and fines. For more on one-time passwords, see Section 28.4.1.

28.3.4 Risk of Weakest Link. One of the frustrations of passwords is that users have to remember too many. Thus, users tend to repeat selection of the same password at multiple sites. This is a serious risk. Exposure of a user password at a poorly maintained site can lead to penetration of the user's account at numerous other sites. It is not easy to deploy technical measures to protect directly against this risk. A particular site can force a user to pick a complex password or can even choose the password for the user. However, it cannot prevent use of the same password elsewhere. This is one area where user education, awareness, and self-interest are paramount.

Malicious attackers can set up rogue Websites easily, to entice users to register for attractive services, whereupon the user's password for other sites may be revealed.

Password safes are available in which difficult-to-guess passwords can be stored and accessed once a master password is supplied. These tools can be helpful, but they increase the importance of creating a suitable master password and of safeguarding it against discovery.

A technical solution to mitigate this problem is to avoid the requirement that a user has to register at multiple sites with user IDs and passwords. Instead, the user should register at a few trusted sites, but the user ID should be usable at multiple sites by secured sharing of assurances that the user has in fact been identified and authenticated sufficiently for business to continue. This is essentially what public key infrastructure (PKI) seeks to do. Once a user has identified him- or herself to a provider of such electronic credentials (client certificates), the certificate becomes a method for authentication, For example, some banks provide a service that allows a user to create a unique number that can substitute for their credit card number in a particular transaction. The user authenticates to the bank site, obtains a unique certificate that substitutes for the credit card number, and gives it to the merchant. The merchant can verify that it generates a valid transaction authorization, but never knows the original credit card number. With authentication based on client certificates, it is not necessary to expose a user's password to multiple sites. An effective marriage of passwords and PKI would reduce the exposure to the weakest link. For more details of PKI, see Chapter 37 in this *Handbook*.

A similar approach stores sensitive information in one place and then directs businesses to that place for payment information. For example, today a number of systems (e.g., PayPal) allow a user to register credit card information once, with a trusted service, and then pay online retailers (e-tailers) via that service.

28.3.5 Risk of Online Guessing. Authentication systems are susceptible to guessing attacks. In *online guessing,* an attacker tries to authenticate using a valid user ID and a guessed password. If the password has been poorly selected, the attacker may get lucky. The attacker also may be able to exploit personal knowledge of the victim to select likely passwords. This approach exploits the documented tendency of naive users to select passwords from lists of obvious words, family, friends, pets, sports, commercial brands, and other easily obtained information. For example, studies of password files consistently show that the most frequently selected password in the world is "password"; the second most frequent is the user ID itself or the user ID backward. An account with the same password as the user ID is often called a *Joe account*, as in User ID: joe; Password: joe.

Another kind of password vulnerable to guessing is a password assigned by default; for example, many software installations create accounts with the same password on all systems. Documentation usually warns users to change those *canonical passwords*, but many people ignore the warning. Canonical passwords are particularly dangerous when they grant access to powerful accounts, such as root accounts or to support functions.

The first line of defense against online attacks is to enforce password complexity rules, in addition to user education and awareness. Many systems today require a minimum of eight-character passwords with a mix of upper- and lower-case letters, numerals, and possibly special characters. Nonetheless, online guessing attacks are still possible, and system logging (see Chapter 53) or application logging (see Chapter 52) can be helpful in identifying successful impersonation. For example, log files may

show that a particular user has never logged on to a particular account outside working hours, yet someone has logged on as that user in the middle of the night.

Some systems react to online attacks by a simple rule that locks the account after a certain number of failed attempts. This rule may have been borrowed from a similar rule with ATM cards. The rule actually makes sense in the context of ATMs, with two-factor authentication based on possession of the card and knowledge of the PIN. However, in a password-only scheme, the "three strikes and out" rule can lead to denial of service to legitimate users. An attacker can easily lock up many accounts by entering three wrong passwords repeatedly, as discussed in Chapter 18 about denial of service in this *Handbook*. A more graceful rule slows down the rate at which password guessing can be attempted, so that a legitimate user may be perceptibly slowed down in authentication but not denied. For example, locking an account for a couple of minutes after three bad passwords suffices to make brute-force guesswork impractical.

In addition, intrusion-detection systems can be configured to alert system administrators immediately upon repeated entry of bad passwords. Human beings then can intervene to determine the cause of the bad passwords—user error or malfeasance.

28.3.6 Risk of Off-Line Dictionary Attacks. The paramount technical attack on password-based authentication systems is the *dictionary attack*. Such attacks start with copying the password file for a target system and placing it on a computer under the attacker's control. The password file normally uses *one-way encryption* that allows the system to encrypt an entered password and compare it to the encrypted form of the legitimate password. If the two encrypted strings match, the entered password is presumably correct, and so the system authenticates the user ID.

The dictionary attack is described as off-line because the attacker obtains the necessary information to carry out the attack and then performs computations off-line to discover the password from this information. It is a guessing attack because the attacker tries different likely passwords from an extensive list of possible passwords (the *dictionary*). The list of likely passwords is called a dictionary because it includes words from one or more natural languages, such as English and Spanish; specialized versions used with *password-cracking* programs may sort words by frequency of use rather than alphabetically to speed up successful guesses.

The initial response to dictionary attacks was to stop users from selecting passwords that could be cracked via a dictionary attack. In essence, the system would try a dictionary attack; if it succeeded, it would prohibit the user from selecting this password. This is not a productive approach because attackers' dictionaries are often ahead of the system's dictionaries. The productive approach is to prevent the attacker from collecting the information necessary to carry out the dictionary attack.

Designers of password-based authentication systems were slow to recognize the risk of dictionary attacks. It has long been understood that passwords should not be stored on a server in cleartext because this becomes a single point of catastrophic failure. Time-sharing systems of the early 1970s stored passwords in a "hashed" form. Knowledge of the hashed form of a password did not reveal the actual password. Authentication of passwords was achieved by computing the hash from the presented password and comparing with the stored hash. The UNIX system actually made the hashed form of user passwords easily readable, since reversing the hash was correctly considered computationally infeasible. However, knowledge of the hashed form of a password is sufficient for dictionary attacks. The attacker guesses a password from a list, or dictionary, of likely passwords, computes its hash, and compares it with the stored hash. If they match, the attacker's guess is verified; otherwise the attacker tries

another guess. Since the late 1980s, UNIX systems have stopped making the file of hashed passwords easy to read, so this vulnerability has been reduced.

UNIX also introduced the concept of a *salt* to make dictionary attacks more difficult. The user password and a random number called the salt are hashed together and stored on the server. The salt itself is also stored on the server. To authenticate a user, the presented password and the stored salt are hashed and compared with the stored hash value. Use of a salt means that a separate dictionary attack is required for every user, since each password guess must be hashed along with the salt. Otherwise, the same attack could be run simultaneously against multiple presented passwords.

28.3.7 Risk of Password Replay. If a password is transmitted in cleartext from client to server, it is susceptible to being picked up on the network by an intruder. This is called *password sniffing*. Many systems require the password to be sent to the server in cleartext. Others require transmission of a hash of the password (usually without a salt). Transmitting the hash of a password is risky for two reasons:

1. The hash is sufficient for a dictionary attack unless a salt is used and kept secret.
2. The attacker does not even need to recover the password. Instead, the attacker can replay the hash of the password when needed.

Many existing systems are susceptible to sniffing of passwords on the network in cleartext or hashed form. Fortunately, technical solutions to this problem do exist.

One approach to the replay threat is to use the server's public key to encrypt any transmission of password-related information to the server: Thus, only the server can decrypt the information by using its private key. This is essentially what server-side Secure Shell (SSH) and Secure Sockets Layer (SSL) do. The server-side mode of both SSL and SSH require the server to have a public key certificate. The client-side mode of these protocols requires that the client also have a public key certificate. This approach can be effective but has its own risks.

An alternate approach is to avoid transmitting the password but instead to employ a protocol that requires knowledge of the password to run successfully. One of the earliest and best-known systems to take this approach is Kerberos. In this system, a user's password is converted to a secret key on the client machine and also stored on the Kerberos server. When the user requests authentication, the Kerberos server sends the user's machine a secret session key encrypted using the shared secret key derived from the user's password. The ability to decrypt this message correctly demonstrates knowledge of the password without actually transmitting it, in cleartext, hashed, or encrypted form. Unfortunately, the Kerberos protocol is susceptible to dictionary attacks; any client machine can pretend to be any user and can obtain the necessary information required for a dictionary attack.

Kerberos also does not use a salt, so the same dictionary attack can be applied to multiple users at one time. Kerberos Version 5 provides for a preauthentication option, which makes it somewhat harder to gather the information for a dictionary attack. The data are no longer available by simply asking the Kerberos server for them; instead they must be sniffed on a network. Recent experiments have shown that dictionary attacks on Kerberos are very practical, so this is a serious vulnerability of a widely deployed password-based authentication system.

Since the early 1990s, many password-based authentication protocols have been published that do not suffer from the dictionary attacks to which Kerberos is so

vulnerable. In particular, *zero-knowledge password proofs* are based on the idea that two parties (computers and people) can demonstrate that they know a secret password without revealing the password. These methods depend on the ability to establish that the two parties both independently selected the same number—but without knowing what the specific number is.[2] One popular conceptual model of this process is the zero-knowledge password proof, which runs as follows:

1. Two people want to test whether they share a secret number (in this thought experiment, a single digit between 1 and 10). In this example, the shared number is 3.
2. The two people have a deck of 10 blank cards.
3. The first person counts down to the third card and makes a mark on the right edge of that card.
4. The deck of cards is arranged so that the second person can mark the left edge of the cards but cannot see the right edge.
5. The second person also counts down to the third card (in this example) and marks the left edge.
6. The card deck is shuffled so that the sequence order is lost and then displayed to both parties.
7. If a single card has a mark on both the right edge and the left edge, then the two parties share the secret number, but neither had to reveal exactly which number it was.[3]

It will be interesting to see if this approach to authentication can be implemented on actual computers and if it is commercially used on a significant scale in the coming years.

28.3.8 Risk of Server Spoofing. As mentioned earlier, one widely used approach to preventing password exposure in transit on a network is to send passwords from client to server encrypted using the server's public key. The server, which has the corresponding private key, can decrypt the password to recover it. Knowledge of the public key is not sufficient for a spoofer to determine the private key. Naïve protocols for protecting the private key can be susceptible to replay attacks. However, there are two well-designed protocols in widespread use today.

Server-side SSL is the protocol that has been used by most Web surfers. In this protocol, the server's public key is used to secure transmission of the user's password from client to server. Like all public key–based schemes, the Achilles' heel of this protocol lies in authentic knowledge of the server's public key. The technology of public key certificates seeks to provide public keys with good assurance of the identity of the server to which they belong. A full discussion of issues with Public Key Infrastructure technology appears in Chapter 37, but it suffices to observe that there are pitfalls with the use of certificates for authentication of servers. A rogue server can collect a user's password by pretending to be something other than what it is. Relying on the look and feel of a Web page for server authentication is hardly sufficient, since it is easy to copy an entire Web page for use as a decoy and to establish confidence before capturing confidential information. An improvement on Website authentication is to associate a specific image and identifying strings (e.g., a picture of a hippopotamus labeled "Archie's Favorite Critter") with a user ID; the Website authenticates itself to

a specific user (e.g., Archie) by displaying that user's chosen image.[4] Authenticity of the server's certificate can be spoofed in many ways that are hard for the user to detect, and manipulation of the trusted root certificates that are configured in the user's Web browser is possible. Moreover, while trust ultimately chains up to a root certificate, the owner of a single certificate below a trusted root is capable of considerable mischief. Server-side SSH is a similar protocol, typically used to provide secure remote access to UNIX servers. Server-side SSL and server-side SSH share the same fundamental vulnerabilities: Both can be spoofed by certificate manipulation. The use of server-side SSL to protect transmission of passwords from client to server is prevalent on the Internet today, but it is important for customers of authentication products to understand the risks inherent in this approach.

In the client-side mode of these protocols, there is no need for a password to be transmitted from client to server, since client-to-server authentication is based on the client's use of its own private key to generate a digital signature. Hence, client-side protocols are not vulnerable to password capture by server spoofing.

28.3.9 Risk of Password Reuse. The need to change passwords with some reasonable frequency is well recognized, but what is reasonable frequency? And how draconian should the enforcement be? It seems that security administrators have pushed too far on these questions. Forcing users to change passwords every month, and enforcing such rules ruthlessly, actually could lead to less security rather than more because so many frustrated users write down and store their ever-changing passwords in nonsecure places. There is a real risk here, created by well-meaning security administrators who have made the problem worse than it inherently is.

Systems that choose a password for the user have their own set of problems and are generally too user unfriendly to be viable in the Internet age. This discussion focuses on systems that allow users to select their own passwords.

How does exposure of a password increase with time? Even the strongest password-based system, with immunity to off-line dictionary attacks and password capture by server spoofing, faces increased exposure as time passes. Over a long period of time, a slow, ongoing, online guessing attack could be successful. Also, the likelihood of inadvertent disclosure by surreptitious observation, or exposure on a Trojan horse–infected computer, increases with time. Nevertheless, a good password, carefully chosen by the user to be safe from dictionary attacks and well memorized by the user, should not be changed casually. A change every six months may be appropriate for well-chosen, brute-force–resistant passwords.

Enforcing password changes is a complicated business, and one where the security community has not really done a good job. It is not difficult to keep track of the age of a password and to force a change when an appropriate time has passed. The difficulty is in forcing the new password to be independent of the old one. In fact, the likelihood is that the new password will be a slight variation of the old one. For example, appending the numeral designating a month to a fixed string enables users to have almost the same password, even if they are forced to change it every month. Some systems will keep a history of recently used passwords to prevent their reuse. A system that keeps a history of, say, five passwords can be fooled by rapidly changing the password six times. To prevent this, there are systems that will not allow a user to change password more than once a day. This has the unfortunate effect of actually increasing risk of password exposure, since a user who realizes that the current password may have been inadvertently exposed cannot change it, exactly when the need to do so is greatest.

28.3.10 Authentication Using Recognition of Symbols. An interesting new approach to user authentication is recognition of particular faces from among a large selection of random photographs. Passfaces software works in this way:

> [A] user sets up an array of photographs and puts some familiar ones into the pool to use as keys—the faces of people the user recognizes; then the software can produce a 3×3 grid of random selections including one of the key pictures. The user picks out the familiar picture and then repeats the exercise twice more with new sets of eight strangers and one friend to authenticate the user.[5]

A white paper explains how human beings are particularly good at recognizing faces; indeed, it seems that we have special circuits that have evolved for rapid and accurate perception of faces. According to the paper, advantages of "using Passfaces over passwords" are that Passfaces:

- Can't be written down or copied
- Can't be given to another person
- Can't be guessed
- Involve cognitive not memory skills
- Can be used as a single or part of a dual form of authentication

The power of the system is enhanced by setting parameters to interfere with misuse of the faces. For example:

> In some high-security applications the grids of faces may be displayed only for a very short time. A half second is long enough for practiced users to recognize their Passfaces. Combined with masking (faces in a grid are overwritten with a common mask face) it is extremely difficult for "shoulder surfers" to learn the Passfaces as the user clicks on them.

28.4 TOKEN-BASED AUTHENTICATION. Token-based authentication relies on something that the user possesses that no other user of the identifier is supposed to possess or be able to access. This authentication can be achieved in many ways, including:

- Access cards
- Proximity and Touch Cards
- Smart cards and dongles
- One-time password generators
- Soft tokens

For details of authentication in physical-security systems, see Chapter 23 in this *Handbook*.

28.4.1 Card Entry Systems.[6] The means of encoding identification data on the cards include optical bar code, magnetic stripe, smart cards with embedded chips that store biometric data, and cards with embedded bits of metal. Most bar codes are not secure; the cards are easily duplicated. Although the newer, two-dimensional bar codes are nearly tamperproof, they cannot store much information. Magnetic-striped cards also have many drawbacks: They cannot store much data, and they are easily

altered, copied, or erased (often by accident). In time, the magnetic data decays and must be reprogrammed. Magnetic card readers are not practical outdoors, because of weather and vandalism. Heavily used magnetic card readers and the cards themselves wear out quickly.

Cards with embedded metal bits are effective. The encoding cannot be seen except by X-ray, and it is durable and permanent. The cards must be held against, or inserted into, a reader that scans the card for the position of each bit. They hold very limited data, the coding is factory installed and cannot be changed, and spare cards must be inventoried. In addition, with the right equipment, these cards can be copied.

Methods of using entry cards include proximity, touch, insertion into a slot that returns the card when it is authenticated, or swiping the card through a narrow channel or in and out of a slot. Swiping is fast, but the card must be hand held. Inserting a card has the same shortcoming as swiping and is slower. Wear, weather, or vandalism can damage the card and card reader.

28.4.2 Proximity and Touch Cards.[7] The best of the new card access control systems use proximity or touch cards. These cards communicate with readers using infrared or microwave transmissions. The reading device powers some types of cards, while others contain miniature batteries. Physically, the cards and card readers are weatherproof, vandal resistant, and do not wear out. Proximity card readers can be surface mounted, recessed flush into a wall, or entirely concealed within a partition so that they do not call attention to a security door.

Touch cards are functionally similar to proximity cards, but they must be held briefly against a reader that is usually visible. Touch cards cost a little less than proximity cards and are good only for entrances with little traffic. The touch-card system is slower, and the cards more easily lost, stolen, or forcefully taken.

Proximity cards (which include RFID cards as well) may be used while concealed inside a pocket, handbag, or wallet. Some are worn concealed or hung on neck lanyards. A proximity card that is also an ID badge that everyone in the workplace wears at all times can access both doors and workstations without being touched. Temporary badges customarily are issued to all visitors, even when escorted, and can be used for access control and to monitor areas entered. Temporary badges are quickly activated with specific privileges and can be revoked automatically and immediately when necessary. Increasingly, the visitor cards are created quickly on site with the visitor's picture and perhaps biometric data as well printed on the card. Longer-term temporary badges can be issued to vendors, contractors, and external employees, although it is best that security personnel store visitor badges safely while the person is off the premises.

The new systems provide many useful functions. They are usually laminated and sealed to prevent wear, damage, or alteration. Individual cards are quickly prepared, activated, and canceled—all on site. The system can restrict entry to specific places, days, and times, and holiday restrictions also can be programmed. Any card can be locked out immediately if lost or stolen, or when the owner leaves.

The newer 13.56-MHz proximity cards function up to three feet away from the card reader; older cards were limited to a range of about four inches. The newer cards are also faster, hold more data, and offer more functionality. Many of the card readers also can write data to the card. There is a trade-off, however, between useful operating range and the amount of data stored. The farther the range, the less the data stored. Most proximity systems are adjustable to optimize distance and speed. For example, on outer perimeter doors, where quick, convenient access is more important than tight security, the systems are set for maximum range. Inner doors that need higher security

are adjusted to utilize more information and to function at a shorter distance, which is still far greater than the older systems allowed. It is not easy without very high-tech equipment, but proximity and touch cards can be compromised.

There are also self-expiring visitor badges that noticeably change color or prominently display the word "expired" after an elapsed period. Self-expiring badges are reusable and come with a fixed expiration period that is usually from two to 24 hours following each activation. These badges cannot be reactivated, except with very sophisticated equipment.

Cards are not the only proximity or touch devices. Keys or patches also are used. The keys can be small, rugged, and easily attached to a key ring or to a small handheld wand that a security guard might use. The patches work in place of touch cards, or with separate access control systems, to upgrade existing legacy systems. The patches are about the size and thickness of a quarter and are easily attached to anything a person normally carries, such as an ID card or badge, a pager or cell phone, or the inside of a wallet. The newer RFID devices will be even smaller.

Card access often is used for all equipment rooms containing servers, network components, or telephone gear; for off-hour access to information systems by users, technicians, and administrators; and for any areas where high-value items are stored. Card systems usually are integrated with premises security to control access to and egress from the building, elevators, service areas, parking, and restrooms, and other parts of the information infrastructure that can also take advantage of the access controls. Plan ahead and consider where additional access control points may someday be needed. Piecemeal additions at a later date can be costly.

Each entry into a controlled area should be logged in a way that cannot be compromised. Logs should provide an accurate audit trail of everyone who sought entry, when, and whether access was denied. Where stronger security is needed, each egress should be logged in the same way. The logging system is best monitored by software that can review all system data in real time, flag trouble quickly, issue periodic summary reports, and quickly search and summarize unusual events. Reviewing logs manually is a cumbersome, time-consuming task. If only manual auditing is possible, there must be a firm policy to do this every few days.

Token-based entry systems by themselves do not provide strong protection. Therefore, some degree of authentication is required.

28.4.3 Smart Cards and Dongles. Another form of token is a smart card. These cards can go into a PC card reader or can be read by a specialized reader. *Dongles* are smart cards that fit into input-output ports such as Universal Serial Bus (USB). A smart card has its own processing capability and typically stores a private key associated with the user. Often, a password or PIN is required to access the card, thereby providing two-factor authentication capability. The smart card enables user authentication by signing some challenge presented to it with the user's private key. The signature is verified by means of the user's public key. A complete discussion of such smart cards involves consideration of public key cryptography, public key certificates or so-called digital certificates, and supporting infrastructure or PKI (see Chapter 37 in this *Handbook*). Suffice it to say that smart cards have long been considered essential for widespread use of public key technology but so far have not been widely deployed.

Hardware tokens offer the potential for stronger authentication than passwords but have seen only limited use due to their perceived costs and their infrastructure requirements. Whether they can be deployed in a scale of millions of users remains to be seen. Authentication by tokens is really authentication by something that the

token knows. Since tokens can be programmed to remember and use secrets much more effectively than humans can, they offer the potential of strong cryptographic authentication, but tamper-proof tokens are not easy to produce. In recent years, attacks based on differential power analysis have proven effective in determining the secret keys and PINs stored on smart cards. These attacks require physical access to a card whose loss probably would be known, so although they may not always be feasible, they certainly call into question the presumed tamper-proof nature of smart cards. As smart cards are more widely deployed, other ingenious attacks are likely to be pursued. Smart cards are more susceptible to secret extraction than tokens because their computations leak such information in the form of electromagnetic radiation.

In comparing tokens with passwords, one can argue that undetected theft is easier with passwords. The token is a physical object whose absence is noticeable. However, tokens create their own problems. Like password generators, they are vulnerable to physical damage and compromise. Smart cards typically are built from thin plastic with a chip embedded in them, making the entire unit susceptible to failure from damage. In addition, users typically may store cards with other credit cards in a wallet or pocket, either of which presents additional environmental hazards. However, for certain industries, the cost of replacing lost or damaged tokens may be worth the reduction in undetected sharing, since a token can be used by only one person at a time.

28.4.4 Soft Tokens. The idea of *soft tokens,* or *software tokens,* has been proposed as a low-cost alternative to hardware tokens. Early soft tokens consisted of a user's private key encrypted with a password and stored on some transportable medium, such as a floppy disk. Such a scheme is extremely vulnerable to dictionary attacks because a guessed password can be verified easily (by testing a putative private key to see if it decrypts a message encrypted using the user's known public key). Moreover, the physical transport of floppy disks and the possible lack of floppy disk drives have led people to store these soft tokens on network servers so they are accessible as needed. Unfortunately, this location also makes them easily accessible to attackers. Protecting access to soft tokens on a network server by means of a password simply returns to the problems of password-based authentication.

It has been suggested that a user's public key could be kept secret and known only to trusted servers to avoid dictionary attacks on the encrypted private key. This approach comes at the severe cost of a closed PKI rather than an open PKI. Schemes for retrieving the private key by means of a secure password-based protocol have been published and are being implemented by some vendors. These schemes ultimately revert to password-based authentication as their foundation. Schemes based on splitting the private key into two parts have been developed. One part of the private key is computed from the password; the other part is stored on an online server, which functions as a network-based virtual smart card. Both parts of the private key are needed for user authentication but are never brought together in one place.

An alternative scheme would be to store the user's entire private key on an online server and make its use contingent on a secure password-based protocol. This approach allows the server to impersonate any user at will and may not be suitable in all environments. However, in all cases, the security of the token relies on the integrity of the computer that uses it. Newer implementations of soft tokens rely on asymmetric cryptography to eliminate the security concerns with storing private keys in a file or other transportable location. The soft token can generate its own key pair and exchange public keys with the authentication server. Although this increases the security of the soft tokens, the concept is inherently weak and prone to attack.

The U.S. Federal Government established a standard for the identity cards it has mandated for all federal employees and contractors. The original Homeland Security Presidential Directive 12 (HSPD 12), dated August 27, 2004, was entitled "Policy for a Common Identification Standard for Federal Employees and Contractors." It defined these requirements for secure and reliable identification that—

- Is issued based on sound criteria for verifying an individual employee's identity
- Is strongly resistant to identity fraud, tampering, counterfeiting, and terrorist exploitation
- Can be rapidly authenticated electronically
- Is issued only by providers whose reliability has been established by an official accreditation process[8]

28.4.5 One-Time Password Generators. A popular form of token, from vendors such as RSA Data Security Inc. and CryptoCard, displays a one-time password, typically a six- or eight-digit numeral, which changes each time an access button is pushed or when a given time has elapsed since the password was last used. The user authenticates by entering the user ID and current value displayed by the token.

The password is called one-time because it expires at the end of its allowable period for use. The token is typically contactless, in that it does not need electrical contact with the computer where the user is presenting authentication data. The user transfers the necessary information from the token via a keyboard or other input device. To make this a two-factor authentication, a fixed user password is also required in addition to the changing one-time password displayed by the token. These tokens are based on shared secret keys, so both the token and the server have a shared secret. The server and the token need to be initialized and then kept synchronized for this scheme to work. In the case of RSA's SecurID, if the time discrepancy between a specific token and the authenticating system exceeds a specified limit, the authenticating software adjusts a value in an internal table to compensate for the time slippage. Vendors such as CryptoCard have developed event-based authentication algorithms to solve the "slippage" problem. These tokens use a seed with the algorithm to generate a unique value for each button push or other activating action. After the next login, that value is now "known" to the authenticating system, and a new value must be provided. The token's value increments based on the seed, without requiring synchronization with the authenticating system.

Password generators must be protected against physical tampering. These devices typically include several measures to cause destruction of the electronic circuits if the outer case is opened; for example, in addition to epoxy-resin glue, tokens may include light-sensitive components that are destroyed immediately by exposure to light, rendering the unit unusable. Some password generators and smart cards cannot even be opened to replace batteries, so the entire card must be replaced on a predictable schedule. Tokens of this kind are available that are guaranteed to last one year, two years, or three years without having the batteries wear out.

28.4.6 Authentication Using Mobile Devices. Another form of token-based authentication increasingly available for access to restricted Websites such as bank account pages registers a mobile device such as a cell phone or a tablet that is capable of receiving Short Message Service (SMS) text messages. During login, the authorized user receives a one-time code to enter in the Web page. This approach puts

additional value on security measures to protect the mobile device; for example, short timeouts requiring entry of a PIN, encryption of data in the device's memory, and methods for remotely inactivating or wiping the device if it is lost or stolen.

28.5 BIOMETRIC AUTHENTICATION. Biometric authentication looks like an excellent solution to the problem of authentication in cyberspace. However, there are several challenges in implementing biometrics, including drawbacks in accuracy (false-positive and false-negative results), loss of biometric identifiers, security of templates, and privacy concerns. For a full treatment of biometric authentication, refer to Chapter 29 in this *Handbook*.

28.6 CROSS-DOMAIN AUTHENTICATION. As more systems and processes become Internet enabled, people come to expect a seamless experience between organizations and applications across the Internet. Furthermore, the ever-increasing risk of identity theft has made individuals and organizations more careful about sending too much identity information across an untrusted network. Over the past several years, efforts have increased to enable easy but secure sharing of authentication and authorization information between organizations. Using the Security Assertion Markup Language (SAML), researchers are developing methods enabling one organization to share just enough information between systems to enable a transaction, without compromising privacy. Shibboleth, a project started in 2000 as an Internet2 middleware initiative, is an open-source project using SAML.[9] Organizations using Shibboleth as a basis for designing new federated authentication and authorization implementations include InCommon[10] and the UK Access Management Federation for Education and Research.[11] To facilitate this kind of information sharing, the participating organizations need to share details about their security policies and procedures, so each may decide in advance whether it will trust the authentication assertions of the other. This kind of information sharing is accomplished most effectively via the policies and practice statements used in a PKI. For details on PKI and certificate requirements for cross-domain authentication, refer to Chapter 37 in this *Handbook*.

28.7 RELATIVE COSTS OF AUTHENTICATION TECHNOLOGIES. One of the frequent responses from security professionals in discussions of identification and authentication using anything other than passwords is that the expense of buying new equipment is prohibitive. However, passwords are not free. In an analysis of the costs of managing passwords, RSA Data Security, makers of the SecurID token, estimated deployment costs (initializing user accounts) at about $12 per user over three years and management costs (replacing forgotten passwords and resetting locked accounts) at about $660 per user over three years. Worse yet, the well-established failure of most users to select strong passwords (i.e., those resistant to guessing, dictionary attacks, and brute-force attacks) makes passwords a weak authentication mode in practice. In comparison, token-based and biometric authentication are more readily affordable and more effective than passwords.[12]

28.8 CONCLUSIONS. Identification and authentication are the foundations for almost all other security objectives in cyberspace. In the past, these problems often were viewed as simple ones whose solution was assumed, before the real security issues came into play. Important standards for computer security were published and practiced without much attention to identification and authentication. In cyberspace,

the problems associated with I&A are severe, and they have barely begun to be solved effectively. A robust, scalable identification and authentication infrastructure is vital to achieving security. Technologies such as tokens and biometrics hold out considerable promise, but their deployment requires infrastructure costs dominated by the cost of hardware for readers and the like. Meanwhile, passwords continue to be strengthened, as we better understand the real risks in using them and as we develop technical means to mitigate those risks. The fact that modern operating systems continue to provide simple password-based authentication, vulnerable to dictionary attacks, reflects poorly on the pace at which security technology is adopted in the marketplace. Looking to the future, one can predict that we will see a mix of passwords, biometrics, and tokens in use, perhaps in two- or three-factor configurations. Biometrics and tokens are likely to dominate the high-assurance end, while passwords will dominate the lower end.

One of the misconceptions that security specialists should seek to dispel is that identification and authentication are *sufficient* for improving public safety. However, assigning a reliable and nonrepudiable identity to someone is in no way equivalent to asserting the trustworthiness of that individual. In closed populations such as employee pools, employers can check the background of potential employees and monitor the performance of existing employees; under those circumstances, knowing someone's identity at the entrance gate may indeed improve security. However, when dealing with a large number of unscreened people, such as potential air passengers, confidently being able to name them tells us nothing about their trustworthiness. Having a clerk at a government office glance at fuel-oil invoices and a birth certificate before granting someone a photo ID is no basis for assuming that the carrier of the valid ID is an inoffensive traveler. As several writers have noted, unambiguously knowing the name of the suicide bomber sitting next to you in a plane is not a reasonable basis for complacency. Timothy McVeigh, the Oklahoma City bomber, was a perfectly identifiable citizen of the United States, but he committed his atrocity nonetheless.[13]

Since September 11, 2001, the air transport industry has been a very public example of both the difficulty of strong identification and authentication and the use of identification and authentication as a public relations substitute for substantive security involving thorough passenger screening. The inherent difficulties of authentication are becoming evident to the lay public and to political and corporate leaders. It is very difficult, perhaps even impossible, to guarantee foolproof identification and authentication in free societies. As technologists, we realize that absolute guarantees cannot be achieved in cyberspace. Too many security professionals seek absolute goals, and too many security technologies are marketed as being stronger than they really are. Our profession will benefit greatly if we address practical problems with practical cost-effective techniques and develop a sound security discipline that contains, bounds, and mitigates the inevitable residual risk that we must face in any large-scale human situation.

28.9 SUMMARY. Passwords are widely used in practice and will continue to be a dominant form of user authentication. There are many risks in deploying passwords, and a number of widely used password systems have serious vulnerabilities. Nonetheless, technical measures can mitigate the inherent vulnerabilities of passwords. Although it takes great skill and care, with our current understanding it is technically possible to build and deploy strong password-based authentication systems using commercial products. The truly inherent risks of undetected theft and undetected sharing can be largely mitigated by new technologies, such as intrusion detection systems. Undetected

sharing may be deterred further by a system that couples high-value secret data, such as credit card account numbers, with passwords. Tokens are available to generate one-time passwords or to communicate directly with authentication systems. Although costs have been dropping, tokens are still not as widely deployed as early predictions suggested they would be. Biometric authentication has been implemented only infrequently and on a small scale but offers great potential, especially for high-security applications. Interesting new research and applications are extending the use of authentication (and authorization) over untrusted networks between federated organizations.

28.10 FURTHER READING

Anderson, J., and R. Vaughn. *A Guide to Understanding Identification and Authentication in Trusted Systems* (1991). "Light Blue Book" in the Rainbow Series, NCSC-TG-017; www.fas.org/irp/nsa/rainbow/tg017.htm

Birch, D. G. W., ed. *Digital Identity Management: Technological, Business and Social Implications*. Aldershot, Hampshire: Gower Publishing, 2007

Integrity Sciences. "Bizcard ZKPP: A Zero-Knowledge Password Proof with Pencil and Paper" (2001). www.integritysciences.com/zkppcard.html

Jablon, D., et al. "Publications on Strong Password Authentication" (2001). www.integritysciences.com/links.html

Jain, L. C., et al., eds. *Intelligent Biometric Techniques in Fingerprint and Face Recognition*. Boca Raton, FL: CRC Press, 1999

Kabay, M. E. "Identification, Authentication and Authorization on the World Wide Web," 1998. www.mekabay.com/infosecmgmt/iaawww.pdf

National Institute of Standards and Technology (NIST). Personal Identity Verification (PIV) of Federal Employees and Contractors, REVISED DRAFT, July 2012. csrc.nist.gov/publications/drafts/fips201-2/draft_nist-fips-201-2.pdf

"One-Time Passwords." FreeBSD Handbook, Chapter 14.5. www.freebsd.org/doc/en_US.ISO8859-l/books/handbook/one-time-passwords.html

Ouf, M. O., R A. El-kammar, and A. E. S. Ahmed. *Implementation and Analysis Requirements of Ultra-Lightweight Radio Frequency Identification Authentication Protocols*. Lap Lambert Academic Publishing, 2012

Radhakrishnan, R. *Identity & Security: A Common Architecture & Framework for SOA and Network Convergence*. London: futuretext, 2007

Smith, R. E. *Authentication: From Passwords to Public Keys*. Reading, MA: Addison-Wesley, 2001

Todorov, D. *Mechanics of User Identification and Authentication: Fundamentals of Identity Management*. Boca Raton, FL: Auerbach, 2007

Tung, B. *Kerberos: A Network Authentication System*. Reading, MA: Addison-Wesley, 1999

Vance, J., "Beyond Passwords: 5 New Ways to Authenticate Users," *Network World*, May 30, 2007. www.networkworld.com/research/2007/060407-multifactor-authentication.html

Wayman, J. L. "Biometric Technology: Testing, Evaluation, Results" (1999); www.engr.sjsu.edu/biometrics/publications_technology.html

Windley, P. *Digital Identity*. Sebastopol, CA: O'Reilly, 2005

Wu, T. "A Real-World Analysis of Kerberos Password Security." Proceedings of the 1999 Network and Distributed System Security Symposium; www.isoc.org/isoc/conferences/ndss/99/proceedings/papers/wu.pdf

28.11 NOTES

1. "No Token Resistance: Citi rolls out Digipass authentication devices to biz clients," *Bank Systems & Technology,* May 25, 2006, www.banktech.com/features/showArticle.jhtml?articleID=188103117
2. J.-J. Quisquater, L. Guillou, and T. Berson, "How to Explain Zero-Knowledge Protocols to Your Children," Proceedings of CRYPTO '89, *Advances in Cryptology,* vol. 435 (1989) pp. 628–631. www.cs.wisc.edu/~mkowalcz/628.pdf
3. R. Wright, "Secret Communication Using a Deck of Cards." Abstract of presentation from DIMACS Research and Education Institute Cryptography and Network Security, July 28–August 15, 1997 (Abstracts of Talks Presented); ftp://dimacs.rutgers.edu/pub/dimacs/TechnicalReports/TechReports/1997/97-80.ps.gz
4. M. E. Kabay, "SiteKey Tries to Counter Phishing," *Network World Security Strategies*, April 3, 2007. www.networkworld.com/newsletters/sec/2007/0402sec1.html
5. M. E. Kabay, "Password Management: Facing the Problem," *Network World Security Strategies,* October 11, 2007. www.networkworld.com/newsletters/sec/2007/1008sec2.html
6. This section (28.4.1) was originally part of Chapter 23 on "Protecting the Physical Information Infrastructure," by Franklin Platt, from the 5th edition of this *Handbook*.
7. This section (28.4.2) was originally part of Chapter 23 on "Protecting the Physical Information Infrastructure," by Franklin Platt, from the 5th edition of this *Handbook*.
8. FIPS 201-1, "Personal Identity Verification (PIV) of Federal Employees and Contractors." Gaithersburg, MD: NIST, 2006, p. iv, http://csrc.nist.gov/publications/fips/fips201-1/FIPS-201-1-chng1.pdf
9. Shibboleth, http://shibboleth.net
10. InCommon, www.incommonfederation.org
11. UK Access Management Federation for Education and Research, http://ukfederation.org/
12. RSA Data Security, "Are Passwords Really Free? A Closer Look at the Hidden Costs of Password Security," 2006. www.rsa.com/go/wpt/wpindex.asp?WPID=3384
13. M. E. Kabay, "Airport Safety," 2005. www.mekabay.com/opinion/airport_safety.pdf

CHAPTER 29

BIOMETRIC AUTHENTICATION

Eric Salveggio, Steven Lovaas, David R. Lease, and Robert Guess

29.1 INTRODUCTION **29·2**

29.2 IMPORTANCE OF IDENTIFICATION AND VERIFICATION **29·2**

29.3 FUNDAMENTALS AND APPLICATIONS **29·2**

29.3.1 Overview and History 29·2
29.3.2 Properties of Biometrics 29·4
29.3.3 Identification, Verification, and Authentication 29·5
29.3.4 Application Areas 29·7
29.3.5 Data Acquisition and Presentation 29·8

29.4 TYPES OF BIOMETRIC TECHNOLOGIES **29·9**

29.4.1 Finger Scan 29·9
29.4.2 Facial Scan/ Recognition 29·11
29.4.3 Hand Geometry Scan 29·12
29.4.4 Iris Scan 29·13
29.4.5 Voice Recognition 29·14
29.4.6 Other Biometric Technologies 29·15

29.5 TYPES OF ERRORS AND SYSTEM METRICS **29·16**

29.5.1 False Accept 29·16
29.5.2 False Reject 29·16
29.5.3 Crossover Error Rate 29·16
29.5.4 Failure to Enroll 29·17

29.6 DISADVANTAGES AND PROBLEMS **29·17**

29.6.1 General Considerations 29·17
29.6.2 Health and Disability Considerations 29·18
29.6.3 Environmental and Cultural Considerations 29·18
29.6.4 Cost Considerations 29·19
29.6.5 Attacks on Biometric Systems 29·19
29.6.6 Privacy Concerns 29·20
29.6.7 Legal Issues 29·22

29.7 RECENT TRENDS IN BIOMETRIC AUTHENTICATION **29·22**

29.7.1 Government Advances in Biometric Authentication 29·22
29.7.2 Face Scanning at Airports and Casinos 29·22
29.7.3 Increased Deployment in the Financial Industry 29·23
29.7.4 Biometrics in the Healthcare Industry 29·23
29.7.5 Increased Deployment of Time and Attendance Systems 29·24

29.8 SUMMARY AND RECOMMENDATIONS **29·24**

29.9 FURTHER READING **29·26**

29.10 NOTES **29·26**

29.1 INTRODUCTION. Once exclusively the purview of law enforcement, intelligence, and national security agencies, biometrics—the automated recognition of people based on their physiological or behavioral characteristics—has finally entered the business mainstream as a method of identification and authentication for access to physical and logical infrastructure. Once a pipe dream, biometric authentication technologies have substantially improved security, convenience, and portability over other commonly used methods of authentication. Adoption rates are still slower than some security professionals have desired; but falling costs, improvements in technologies, increased security needs, and changing government regulations are encouraging the adoption of biometrics.

29.2 IMPORTANCE OF IDENTIFICATION AND VERIFICATION. Ensuring the identity and authenticity of persons is a prerequisite to security and efficiency in present-day organizational operations. Intruders can damage physical and logical infrastructure, steal proprietary information, compromise competitive assets, and threaten organizational sustainability. Traditional methods of recognition and identification, wherein one individual identifies another based on his or her voice, physical appearance, or gait, are impractical, inefficient, and inaccurate in the scope of contemporary organizational operations. To address the need for rapid, efficient, and cost-effective authentication, organizations today primarily rely on the two methods of "something you know" and "something you have" (either applied individually or in combination) to verify the identity of persons accessing their physical and/or logical infrastructure. The most robust authentication systems use multiple factors of authentication. These forms of authentication are described in Chapter 28 in this *Handbook*.

29.3 FUNDAMENTALS AND APPLICATIONS. Biometrics are based on the measurement and matching of distinctive physiological and/or behavioral characteristics. The former are based on direct measurement of a physiological characteristic of some part of the human body. Examples of physiological biometrics include finger, hand, retina, face, and iris scans. The latter indirectly measure characteristics of the human body based on measurements and data derived from an action. Commonly used behavioral biometrics include voice and signature scan and keystroke pattern.

29.3.1 Overview and History. The use of nonautomated biometrics dates back to the beginning of human civilization, when individuals first began identifying other individuals based on certain physical or behavioral characteristics. The concept of biometrics as a means of authentication dates back more than 2,000 years. As early as 300 BC, Assyrian potters used their thumbprint as an early form of brand identity for their merchandise. The use of handwritten signatures (chops) in classical China is another example of an early biometric. In the first instance of a formal, legal biometric authentication system, fingerprints were used to sign contracts during the Tang dynasty (AD 618–906).

The development of contemporary biometric systems can be viewed as an outgrowth of the efforts of forensic scientists and law enforcement agencies to identify and classify criminals in the late nineteenth and early twentieth centuries. In 1882, Alphonse Bertillon introduced a system of body measurements called *anthropometry* to identify criminals. This system was proven unreliable because the measurements taken were not globally unique. A student of Bertillon, Edmund Locard, later proposed a fingerprint system based on the work of Sir Edmond Galton to identify people by analyzing unique points in fingerprint ridges and pores. Locard's system, which used 12 Galton points,

is considered reliable to this day. This methodology underlies fully automated modern biometrics, including the Integrated Automated Fingerprint Identification Systems (IAFIS) used by law enforcement agencies. Commercial biometric systems (typically relying on hand geometry) designed for use in physical access to buildings, emerged in the 1960s and 1970s.

Biometric methods of identification and identity verification, including automatic fingerprint analysis and facial recognition technologies, have been available and used by some government/public agencies (e.g., law enforcement, intelligence, and national security) and a few private industries (e.g., facial recognition scans in casinos) since the 1960s and 1970s. Notwithstanding the potential benefits and advantages over other authentication methods, biometrics have not been widely applied, particularly in the corporate world. Analysts cite high costs of equipment and implementation, technological problems, vulnerabilities of specific biometrics, lack of standards, and user resistance (notably, concerns over privacy) as reasons for the lack of implementation.

However, beginning in the 1990s, significant improvements in biometric technologies, a movement toward standardization, regulatory changes requiring organizations to adopt stringent security and privacy controls, and significantly reduced costs have encouraged wider adoption. A number of U.S. Government agencies (e.g., Department of Homeland Security, Department of Transportation, Department of Defense, Customs and Border Protection, Department of Justice, National Library of Medicine) and businesses in certain industries (e.g., health care and finance) have significantly increased their use of biometrics during the past few years—a factor that is likely to encourage other organizations to adopt biometrics as well. HP and Lenovo offer fingerprint scanners on their notebook computers, enabling users to increase security by requiring a finger swipe and a password (or just a finger swipe) to access files. Other computer manufacturers and peripheral vendors have added fingerprint scanning to computer keyboards and/or developed standalone fingerprint scanners that connect to computers via a Universal Serial Bus (USB) port.

Some analysts see the 2001 attacks on the World Trade Center and the Pentagon also as a key impetus behind the increased usage of biometrics for authentication and identification. The terrorist attacks have been critical in encouraging adoption not only because they have heightened the security concerns of companies and agencies but also because the impact and implications of the terrorists' attacks seem to have lowered users' resistance to the use of biometrics by employers and government. In other words, in the same way that the threat of global terrorism reduced public objections to possible infringements on civil liberties as a result of implementation of the USA PATRIOT Act and other security-focused measures, these attacks also appear to have rendered many people less sensitive to the potential privacy-invading implications of biometrics. Boroshok reported in 2007 that a survey sponsored by AuthenTec found that 71 percent of U.S. consumers would pay more for biometric security options in their cell phones and 63 percent of consumers would pay an additional cost for these options to be added to their personal computers.[1]

The changing security environment has prompted forecasts of rapid growth in biometrics. In 2004, the International Biometric Group (IBG) had predicted rapid growth for the biometrics industry over the next several years from revenues totaling under $50 million in 2004 to revenues of almost $200 million in 2008.[2] In late 2003, analysts at the San Jose, California–based market research firm Frost and Sullivan predicted that biometric applications from commercial applications (not including the Federal Bureau of Investigation's IAFIS) would jump from $93.4 million in 2001 to $2.05 billion by 2006—up from the $700 million (in 2006) that these analysts predicted prior

to the September 11, 2001, attacks. In January 2009, IBG released their most recent projections of market growth for the biometrics industry. They estimated that biometric industry revenues will grow from $3.01 million in 2007 to over $7.4 million by 2012.[3] This is a far cry from their initial projected growth rate published in 2004, and may be attributed to a lack of standards, regulatory issues, and privacy groups that decried the invasion of their privacy by these devices (IBG hasn't publicly released a more recent report so we can't tell if these projections were met).

Despite rosy projections from industry analysts, adoption of biometric authentication systems continues to lag. Some firms continue to cite cost issues (though this is dwindling) and privacy concerns, while others point to problems surrounding biometric implementation in airports and among government agencies. Overall, surveys of companies indicate that forecasts of dramatic and rapid growth in biometrics implementation may be overstated.

This has been proven to be true by the recent 2013 Spending Priorities Survey put out by *InformationWeek.* Of the 513 businesses that responded to the survey, 13 percent had budget decreases from their 2012 budgets, with 43 percent having no change in their budgets—for many, it's a "do what you can with what you have" mentality. However, in spite of this, 58 percent of those surveyed wanted to improve their security overall.[4] In another 2013 survey, however, the *InformationWeek* survey "Strategic Security Survey" revealed 52 percent of the more than 1,000 respondents wanted to increase their identity and password management practices.[5] The seeming disconnect may be attributed to management wanting to keep costs down, and the ever-changing world of regulations—many are waiting to see what the landscape will look like when the dust settles.

A 2003 Forrester Research survey found that only 1 percent of companies had implemented biometric systems, just 3 percent had a biometric system rollout in progress, only 15 percent were testing biometrics, and 58 percent of those surveyed had *no plans* to try biometrics. Ten years later, we still have yet to see any substantial adoption of biometric authentication.

29.3.2 Properties of Biometrics. The contemporary meaning of biometrics emphasizes its automated aspects, which allow for deployment on a large scale. The most widely cited definition of biometrics is some variation of "the automatic identification of a person based on his or her physiological or behavioral characteristics."[6] The term "biometrics" generally is used as a noun to refer to the automatic recognition of persons based on their physical or behavioral characteristics. The term "biometric" can be used as a noun in reference to a single technology or measure (e.g., finger scan is a commonly used biometric) or as an adjective, as in "a biometric system uses integrated hardware and software to conduct identification or verification."[7]

Biometrics have long been touted as a possible solution to the problems and vulnerabilities of other commonly used methods of authentication and identification. They represent sophisticated versions of the traditional means of identification, such as a guard allowing access to a user whom the guard recognizes by sight. Biometrics commonly are defined as automated methods of recognition/verification/identification of individuals based on some measurable physiological or behavioral characteristics, such as fingerprints, hand geometry, facial shape, iris pattern, voice, signature, and the like.

Whereas identification (ID) badges and keys authenticate the user based on something the user possesses, and passwords/personal identification numbers (PINs) authenticate the user based on what the user knows, biometrics allows authentication and identity verification based on who the user *is.* Because biometric methodologies of authentication actually base identification on physiological or behavioral "pieces" of the

user, biometrics represents the only form of authentication that *directly authenticates the user*.

Biometrics have a number of other obvious advantages over other commonly used authentication methods. Unlike an ID badge or a USB key, one cannot easily lose or misplace a fingerprint or other biometric measures. Likewise, unlike the case with passwords and PINs, one does not need to remember and one is not subject to forgetting a physiological or behavioral characteristic. Although biometric measures *can* be compromised, in general, a biometric is much more difficult to manipulate by stealing, forging, sharing, or destroying than other commonly used authentication tools. Biometrics also provide considerable convenience, as opposed to the hassle of memorizing dozens of passwords.

Although the initial costs are quite high, the implementation of biometric systems typically results in much lower administrative costs than other access methodologies due to fewer calls to the helpdesk for technical support to reset passwords, no need to issue replacement ID badges, and so on. For these and other reasons, biometrics are viewed as providing better security, increased efficiency, and more reliable identity assurance than other commonly used methods of authentication/identification based on what a user possesses or what a user knows.

In theory, almost any human physiological and/or behavioral characteristic can be used as a biometric measure. However, to fit within a viable, potentially accurate, and practical biometric system, the biometric used should also satisfy four other requirements offered by Jain[8] and Bolle, Connell, Pankanti, Ratha, and Senior[9]:

1. **Universality**. Every person should have the biometric characteristic.
2. **Uniqueness**. No two persons should be the same in terms of the biometric characteristic. Jain proposed the somewhat lower standard of distinctiveness, defined as "any two persons would be sufficiently different in terms of the characteristic."[10]
3. **Permanence**. The biometric should be relatively invariant over a significant period of time.
4. **Collectability**. The biometric characteristic should lend itself to quantitative measurement in a practical manner.

Bolle, Connell, Pankanti, Ratha, and Senior argued that the biometric should also have a fifth attribute: acceptability, defined as "the particular user population and the public in general should have no strong objections to the measuring/collection of the biometric."[11] Jain argued that a practical biometric system should consider two other attributes: (1) performance, which is "the achievable recognition accuracy and speed, the resources required to achieve the desired performance, as well as the operational and environmental factors that affect the performance," and (2) circumvention, which "reflects how easily the system can be fooled using fraudulent methods."[12] Along with this line of thinking, we find that the BYOD explosion among the workplace is opening an avenue for identity and access management. With the acquisition of AuthenTec by Apple in 2012, suppositions were running amok with the proposition of a new mobile-enabled biometric device for use inside and outside the enterprise. While the BYOD era has brought about the "acceptability" attribute, Jain's idea of the proper operational and environmental factors are still there—how largely this will be accepted will be determined by how well these can be managed and centralized using IAM systems.[13]

29.3.3 Identification, Verification, and Authentication. Biometrics can be used to fill several roles related to identification and authentication (for more

background on those concepts, see Chapter 28 in this *Handbook*). These roles are often conflated in the use of the term *biometric authentication*, when in fact the system using the biometric data may be performing identification, verification, or authentication… or some combination of these. To avoid confusion in the use of terms, it is important to distinguish the various ways in which biometrics are used.

Identification systems answer the question, *Who do you claim to be?* They allow for selecting a single subject from a possible group of subjects based on the presentation of the identifying information. Identification systems often are referred to as *1: N* (one-to-*N* or one-to-many) systems because the subject's biometric information is compared against multiple (*N*) records in the attempt to determine which one is a match. This kind of system has application in searching for a positive match (such as the facial recognition one might want to do in airports or stadiums for finding terrorists), as well as a negative match (designed to ensure that a person's biometric information is *not* present in the database, such as preventing people from enrolling more than once in large-scale benefits programs). But biometric identification also can serve as a convenient substitute to entering a username or swiping a card to proffer a claim of identity to an automated entry or logical access control system.

Verification systems answer the question, *Are you who you claim to be?* They attempt to match proffered information (passwords, token identifiers, or biometric data) with previously registered ("enrolled") credentials stored in the system; the match to be attempted is selected by a specific user identifier (ID) presented by the subject. Biometric verification systems are referred to as 1:1 (one-to-one) systems because, while they may contain thousands or even millions of biometric records, they are "always predicated on a user's biometric data being matched against only his or her own enrolled biometric data."[14] Thus, the presentation of biometric information is intended to confirm (or authenticate) the veracity of the subject's claim to the presented identity.

Whether one of these biometric systems qualifies as an *authentication* system depends on the implementation. A biometric verification system, which pairs a user ID with biometric information, always performs both identification and authentication. In that case, the biometric data is being used in place of a known password or the output of an authentication token, and may increase security by avoiding some of the vulnerabilities that plague other authentication credentials. A biometric identification system, on the other hand, does not always perform authentication. It *may* also fill the role of an authentication system if additional authentication factors (such as a PIN) are presented to confirm that the biometric information is in fact being presented by the subject the system identifies as being linked to the information. Without a second piece of information to confirm the linking, the use of a biometric identifier substitutes for the presentation of other user identifiers (such as typing in a user ID), and in this sense is not performing authentication.

It may be argued that because biometric data refers directly to some part of a subject's body or to an observable characteristic of the subject's behavior, presentation of biometric credentials should be considered a form of direct authentication. On this line of argument, in the same way that a security guard recognizes authorized employees at the gate in a combination identification/authentication action, a biometric identification system is in a sense authenticating the subject. This argument, however, glosses over the vulnerabilities that can plaque both approaches: the guard can be fooled through disguises, can misidentify siblings (especially twins), or can simply miss a sufficiently similar face. Similarly, the automated biometric identification system can be fooled by spoofed credentials, as well as failing to distinguish between very similar biometric data. Relying on the single biometric presentation for both identification and

implicit authentication is certainly quicker and more convenient, but may introduce risk that should be mitigated by a second authentication factor.

Biometric identification systems are more difficult to design and implement than verification systems because of the extensive biometric database search capabilities needed. Additionally, identification systems are more subject to error than verification systems, because many more matches must be conducted, matches that increase the opportunity for error. Verification systems are overall much faster (often rendering a match/no match decision within less than a second) and more accurate than identification systems. Verification systems, as opposed to identification systems, predominate in private sector applications, particularly for computer and network security applications. Verification systems also predominate in applications designed to authenticate rights-to-access to buildings and rooms, although sometimes identification systems are also deployed in high-security environments. Identification systems are often found in public sector applications, such as law enforcement (e.g., parole and prison administration, forensics, etc.), large-scale public benefits programs, intelligence, and national security applications.

29.3.4 Application Areas. Although there are many potential applications for biometrics, the primary ones can be divided into four categories: systems security (logical access systems), facilities access (physical access systems), ensuring the uniqueness of individuals, and public identification systems. The common thread among these four applications is that they all rely on individuals enrolled in the systems. The significance of whether individuals are enrolled or not enrolled will be explained in Section 29.3.5.

29.3.4.1 Security (Logical Access Systems). Logical access systems "monitor, restrict, or grant access to data or information."[15] Examples include accessing a computer or network or accessing an account. In these systems, biometrics replace or complement PINs, passwords, and tokens. The volume and value of electronic commerce plus the value of sensitive and personal information transported and/or stored on networks and computers make the use of biometrics to secure logical access a much more robust industry segment than physical security.

The use of biometric technologies for logical access control is still very much in its infancy. The most common biometric approach is to use fingerprint readers, either with a stand-alone USB reader or with a reader embedded in a laptop. Manufacturers are beginning to incorporate the Trusted Platform Module (TPM) chip in new laptops to support a variety of cryptographic applications. In combination with biometric devices like fingerprint readers, the TPM chip can allow applications like Microsoft's BitLocker to apply biometric access control to encrypted volumes on a hard drive. Such technologies are still new enough to suffer problems with backward compatibility and inter-vendor support, but they offer the promise of much more secure logical access.

29.3.4.2 Facilities Access (Physical Access Systems). Facilities access systems "monitor, restrict, or grant movement of a person or object into or out of a specific area."[16] In these systems, biometrics replace or complement keys, access cards, or security cards, allowing authorized users access to rooms, vaults, and other secure areas. Physical access systems often are deployed in major public infrastructure settings, such as airports, security checkpoints, and border facilities, in order to monitor and restrict movements of unauthorized or suspicious persons. In addition to entry to secure rooms, physical access systems, when applied in business settings, include

time-and-attendance systems by combining access to a location with an audit of when the authentication occurred.

Biometric technologies have been in use for physical access control for some time but still represent a range of possible implementations, from stand-alone fingerprint-reading door locks to complete systems with central storage of biometric templates, logging, and power failure protection. Selecting a system from within such a wide product range must fit into the overall security stance of the facility, particularly in regard to the storage of templates. In a simple stand-alone door system, for example, fingerprint templates would be stored in or near the locking device, and the system might not have any logging features. Unless coupled with surveillance and intrusion alarms anticipating physical compromise, these devices are more appropriate for lower security applications, such as storage facilities where physical key access might not be preventing theft. More centrally connected systems, however, offer greater integration into overall monitoring and control of access but suffer from communication and power issues as they push templates across a network. Selecting a product that interacts well with other parts of the overall access control system is crucial.

29.3.4.3 Ensuring the Uniqueness of Individuals. Uniqueness biometric identification systems typically focus on preventing double enrollment in programs or applications, such as a social benefits program. The main use of this application occurs in the public sector, although similar systems could be implemented to prevent double enrollment in employee benefits programs.

29.3.4.4 Public Identification Systems. A final biometric application of note is its use to identify criminals and/or terrorists. Criminals and terrorists can wear disguises, acquire fake documents, and change their names, but biometric data are fairly difficult to forge. In 2004, the United States Department of Defense started the Automated Biometric Identification System (ABIS), which collects biometric data on Iraqi insurgents in a manner compatible with IAFIS. This allows for identification of known repeat offenders and wanted persons. Soldiers also use the Biometric Automated Toolset (BAT) developed by the Army's Language Technology Office to identify persons on the scene of bombing attacks. Anyone present in the area can be cross-referenced with an existing database of insurgents. The BAT is also used to enroll and identify members of the Iraqi army. Although biometric systems hold much promise, it is also important to understand there are still limitations of current technology, but with the advent of nanotechnology, many of these are quickly being overcome.

29.3.5 Data Acquisition and Presentation. As Jain explains, "A biometric system is essentially a pattern recognition system that operates by acquiring biometric data from an individual, extracting a feature set from the acquired data, and comparing this feature set against the template set in the database."[17] The starting point for the biometric system is *enrollment*: A user's biometric data are initially collected and processed into a template, the form in which they are then stored for ongoing use. As Woodward, Orlans, and Higgins explain, "Templates are not raw data or the scanned images of a biometric sample, but rather they are an accumulation of the distinctive features extracted by the biometric system."[18] Liu and Silverman describe the template as "a mathematical representation of biometric data. A template can vary in size from 9 bytes for hand geometry to several thousand bytes for facial recognition."[19] Templates are proprietary to each vendor and technology with little or no interoperability between systems. This lack of interoperability is attractive from a privacy perspective

but unattractive from the perspective of cost effectiveness and the prospective implementer who is concerned about committing significant investment to a single nonstandardized technology.

The term "presentation" refers to the process by which a user provides biometric data to an acquisition device by looking in the direction of a camera, placing a finger on a pad or sensor, or some other specified physiological exam. For purposes of verification or identification, the user presents biometric data, which are then processed and converted to a template. This template is an extraction of distinctive features and is not adequate for the reconstruction of the original biometric data. The scanned template is then matched against the stored enrollment template(s). Each time a user makes a presentation, a new template is created and matched. It is important to note, especially from the perspective of privacy concerns, that biometric systems do not store raw biometric data; instead they use the data for template creation and, in most cases, discard the biometric data. The biometric system's match/no-match decisions are based on a score, which is "a number indicating the degree of similarity or correlation resulting from the comparison of enrollment and verification templates."[20] Like the templates, the scoring system is based on proprietary algorithms; there is no standard system.

29.4 TYPES OF BIOMETRIC TECHNOLOGIES. As previously noted, biometrics generally can be grouped into two categories: physiological and behavioral. The International Biometric Group provides data on comparative market share of various biometric technologies. The IBG data focus on market share from commercial and government applications. The top eight biometric technologies are all from the physiological category (see Exhibit 29.1).

Five of these biometrics are discussed in detail in the next sections.

29.4.1 Finger Scan. Finger scan or fingerprint technology is by far the most widely deployed biometric technology. Finger scan's number-one status as a biometric is maintained even if the extensive use of fingerprinting by law enforcement agencies is excluded. The type of fingerprinting employed in commercial biometric systems differs from the one used in law enforcement. In most commercially available biometric applications, the station provides only for the scan of a single finger on one hand, whereas

EXHIBIT 29.1 Comparative Market Share of Biometric Technologies by 2015

	Market Share
Biometric Technology	**2015 (projected)**
Finger scan	15%
Facial scan	15%
Hand scan	1%
Iris scan	16%
Vein	10%
Voice	13%
Signature	10%
AFIS/Livescan	16%

Sources: (1) Market share data from PRWeb, Retrieved May 29, 2013, from http://www.prweb.com/releases/2007/04/prweb516626.htm (2) J. McHale, "Biometrics: The Body's Keys," *Military & Aerospace Electronics* 14, No. 12 (2003): 17–23.

law enforcement agencies often rely on full sets of fingerprints. In addition to being the most widely used biometric, fingerprinting is also one of the oldest and most well-researched biometric technologies. Because it is a widely used, well-documented, and mature technology, costs for the deployment of finger-scan-based technologies are relatively low. Single-quantity pricing for a workstation version with associated software can be as low as $150; server versions are currently priced as low as $50 per unit.

The strengths of finger scan are one of the principal reasons for its popularity and include:

- Wide use.
- Mature technology.
- Low cost.
- High ease of use. (Very little training is required to place a finger on a finger pad.)
- Ergonomic design. (Comfortable to use for most users.)
- Low error incidence. (False match rates are extremely low; crossover error rate is lower than voice scan and facial recognition, higher than hand geometry and iris scan.)
- Fast transaction times. (In most systems, authentication takes less than a second.)
- Capacity to be deployed in a wide range of environments (e.g., on workstations, doorways, indoors/outdoors).
- Ability to increase accuracy levels by enrolling multiple fingers.
- Can provide identification with a high level of accuracy (if properly configured to include multiple enrolled fingers) in addition to verification.

Despite its multiple strengths, finger scan is not without significant weaknesses. As Chirillo and Blaul note, some of this technology's weaknesses stem from the same factors that lend it its strengths. "Because fingerprint technology is one of the oldest and most well-known technologies, a good amount of information is publicly available on how to defeat it."[21] A number of ways exist to foil finger scans and produce a false match (false accept), including the use of a dummy finger constructed of latex or other material, manipulation of the scanner so as to raise the latent print of the person who used the scanner previously, and even use of an actual finger that is no longer attached to a body. (Most finger scanners cannot discriminate between live and dead tissue.) Because of these factors, the security levels of finger scans are not actually as impressive as the low error rates seem to indicate. It should be noted that countermeasures could be taken to overcome the vulnerability of finger scans to fraud. For example, enrolling additional fingers makes fraud more difficult. To reduce the chance that the system will be foiled by synthetic or dismembered fingers, thermal and/or moisture scanners can be added to the sensors to detect finger temperature and moisture levels that would indicate the vitality of the finger.

Other weaknesses include:

- A scanner requires frequent maintenance because screens/sensors tend to retain an obstructing buildup of user skin oil and residue.
- Performance can deteriorate over time, both because of aging of the users (and wearing away of fingertips) and because of the need for system maintenance.
- Finger-scan biometrics are obviously not appropriate for users with missing hands or hand disabilities.

- Performance levels deteriorate among users who have hand tremors because the presentation of biometric data will be distorted.
- Performance levels also deteriorate when users' fingers are either overly dry (a certain amount of normal skin moisture is needed for an accurate reading) or overly moist or oily (as from too much hand lotion).
- There is a small but significant failure to enroll (FTE) rate even among a population with hands and without disabilities. The FTE rate for finger scans is estimated at 2 to 10 percent and is attributed to persons with genetically indistinct prints, scarred fingers, dry skin, and fingerprints worn down by age and/or manual labor.

Perhaps the biggest weakness of finger scan, however, has nothing to do with the accuracy and reliability of the technology. Instead, it relates to user acceptance. Because of the association of finger scans with law enforcement and criminality, often such scans are not readily accepted by users who dislike the technology's "taint" with forensic applications and who may worry that finger-scan biometric data will be used for other purposes.

According to Chirillo and Blaul, "Another reason fingerprint technology is not highly accepted is that it may require individuals to share or touch the same device that others touch."[22]

29.4.2 Facial Scan/Recognition. Bolle, Connell, Pankanti, Ratha, and Senior note that "face appearance is a particularly compelling biometric because it is one used every day by nearly everyone as the primary means for recognizing other humans. Because of its naturalness, face recognition is more acceptable than other biometrics."[23] However, user acceptance of facial scans drops significantly when users discover that it has been used covertly. As Imparato observes, "Of all the biometric technologies currently in use, face recognition is arguably the most controversial."[23, 24] Like finger scans, face recognition relies on the identification of unknown face images by comparison to a database of (known) face templates. Face recognition is used overtly in access control where it is used for one-to-one identification. This application yields relatively high performance because the environment is highly controlled and input data are predictable. Face recognition may also be used covertly in surveillance where it is used to locate people in crowds (one to many), albeit with mixed results. For example, the city of Virginia Beach, Virginia, has employed a face recognition system to identify known (pre-enrolled) felons for over five years, but has yet to identify a single criminal. A variety of facial recognition technologies, ranging from single image, video sequence, three-dimensional image, near infrared, to facial thermograms, are available.

Facial recognition offers these benefits:

- It has the capacity to leverage existing image acquisition equipment, such as digital cameras, Web, video, and the like.
- Because facial recognition is a software-based technology, it is often unnecessary to purchase new hardware, especially given the number of Closed Circuit Television (CCTV) and surveillance cameras in broad use.
- The lack of need for specialized hardware can help keep the cost of this technology down, assuming that high software costs do not counterbalance the savings from the hardware.
- It is the only biometric capable of identification at a distance without the subject's cooperation or even awareness.

- It is easy to use. All that is required is that the user (or target) look at the camera.
- It does *not* require the user to touch any device (a major objection for some users with finger scans and hand scans).
- When deployed in verification situations, facial scans have extremely low failure-to-enroll rates. (Unlike fingerprints, human faces are almost always distinctive.)
- They are capable of enrolling static images (e.g., photographs on driver's licenses), a factor that makes it possible to implement very large scale enrollments at a relatively low cost and in a brief amount of time.

Facial recognition systems have a number of serious weaknesses too. The predominant weakness (which derives from a combination of the technology's other weaknesses) is the low accuracy and high error rate of this biometric. Whether deployed covertly or overtly, facial recognition has the lowest accuracy rate among all five top biometrics. Tilted heads and low camera resolution still confounds the programs, as evidenced in the Boston bombings of 2013.[25]

Evidence of the technology's low accuracy rate comes from a study at Palm Beach (Florida) International Airport that showed that the system failed more than 50 percent of the time to match the 15 employees who had enrolled in the database for a trial run. Out of 958 pass-throughs, the system matched the employees' faces just 455 times. Some studies suggest that accuracy improvements can be made in facial recognition systems, but these improvements will come at a very high cost. For example, a facial recognition software package from Visionics FaceIt resulted in impressively low error rates, as long as lighting conditions were perfect. The software costs $30,000 for a three-camera system.

Other weaknesses include:

- False matches (false accepts) routinely occur in the case of twins, and most systems are insensitive enough for someone skillful at disguise and impersonation to trick the system into a false match.
- More likely than false matches, however, are false nonmatches (false rejects), which can occur as a result of facial expressions; changes in hairstyle, makeup, facial hair, significant changes in body weight, eyeglasses, and age-related facial changes.
- The acquisition environment can have a dramatic impact on facial recognition system accuracy. In particular, lighting, either too bright or too dim, can dramatically increase the error rate.
- The perceived threat to privacy. Overtly deployed facial recognition technologies (e.g., used for identification and access) are generally judged relatively unobtrusive and meet with a high level of user acceptance. However, covertly deployed systems, such as those used for surveillance, pose significant threats to privacy. This threat is generally viewed as much more serious than that posed by the other top biometrics.

29.4.3 Hand Geometry Scan. Hand geometry scans refer not to handprints or to any analogy of fingerprints but rather to the geometric structure (or geometric invariants) of the human hand. Nanavati, Thieme, and Nanavati explain that "hand-scan technology utilizes the distinctive aspects of the hand—in particular, the height and width of the back of the hand and fingers—to verify the identity of individuals."[26] The

leading hardware maker for this technology, Recognition Systems, Inc. (RSI), has a basic hand scanner that takes upward of 90 measurements from three to four enrollments to create a user template that includes length, width, and thickness, plus surface area of the hand and fingers. Newer systems include temperature-sensing mechanisms to ensure "live" subjects. All the components of a hand scan system (acquisition hardware, matching software, storage components) reside within a stand-alone device. Hand scans are a well-established biometric technology (they have been in widespread use since the 1970s), but compared to other leading biometrics, hand scans tend to be much more limited in their range of applications. Hand scans are used exclusively for verification rather than for identification because the hand measurements are not distinctive or specific enough to allow for identification applications. For this reason, hand scans are used mostly for physical access and time-and-attendance applications. In the latter case, they are used as a way to eliminate the problem of "buddy-punching" whereby one employee punches in or out for a coworker who is not present.

Hand scan technology has changed very little since it was first introduced over 30 years ago, so its strengths and weaknesses are well established. The principal strengths of the hand scan include:

- Operates in very challenging environments. (The equipment is typically unaffected by light, dust, moisture, or temperature.)
- Established and reliable technology.
- Ease of use. (Users simply stick their hand in the unit; placement matters little.)
- Resistance to fraud compared to other biometrics. (It would be difficult and time consuming to substitute a fake sample.)
- Small template size (as low as 9 bytes; much smaller than other biometrics, allowing for storage of thousands of templates in a single unit).
- Based on a relatively stable physiological characteristic.
- High level of user acceptance and lack of attached stigma.

Problems reported in using hand scans include:

- Limited accuracy (which in turn limits its use to verification not identification). The relatively low accuracy of hand scan (higher than facial recognition and behavioral biometrics but lower than finger and iris scans) is a result of the general lack of physical variety expressed in the hand as well as the relatively small number of features measured by hand scan.
- Comparatively large form factor (This limits the technology's deployment in computer-oriented applications that require hardware with a smaller footprint.)
- Some people resent forced contact with possibly unclean surfaces.
- Ergonomic design limits its use by some populations (e.g., the disabled).
- Comparatively high cost. While the cost of palm scanners has come down to around $500 for a USB model, enterprise level scanners are still a rather pricey unit, even with the increase in companies producing them.

29.4.4 Iris Scan. Iris scan technology uses the unique pattern formed by the iris—the colored part of the eye bounded by the pupil and the sclera—to identify or verify the identity of individuals. The iris pattern is unique, for even in the same

individual, no two irises are alike. The uniqueness of iris patterns has been likened to that of multilayered snowflakes. To put this into perspective, the chance of mistaking one iris pattern for another is 1 in 10^{78}! Iris scans accomplish this by allowing for more than 200 points of reference for comparison, as opposed to the 60 to 70 used in fingerprints. The unique aspects of the iris make it an ideal biometric for high-security applications; enrolling both irises from the same individual can enhance the level of security. In addition to high-security physical access applications, iris-scan technology has been used in automated teller machines (ATMs) and banking kiosks. It is also now being used among police departments in handheld units to identify suspects, and is being introduced for such jobs as border crossings, and is being regularly used in India, Iraq, and Dubai.[27]

The most important strength of iris biometrics is its accuracy, the most critical weakness of facial scanning. Of all the leading biometrics, iris technology has the lowest error rate and the highest level of overall accuracy. Other strengths of this biometric include:

- Ability to be used both for verification and for identification.
- Stability of its biometric characteristics over a lifetime.
- Relative difficulty to fake or spoof because it is an internal biometric.
- The fact that the iris is minimally subject to outside influences when compared to biometrics like fingerprints and faces.

The major weaknesses of the iris biometric concern user perceptions and problems in the user-technology interface. Other weaknesses include:

- Acquisition of the image requires moderate training and attentiveness: Users must stand still and look straight into the scanner with eyes open and unblinking.
- Users often report some physical discomfort with the use of eye-based technology, although less so than with retina scanning technology.
- Anecdotal reports also suggest a fairly high level of user psychological resistance to iris-scanning technology, with some users believing that the scanner will lead to eye damage.
- Can be adversely affected by lighting and other environmental conditions (although not to the extent of facial scanning).
- In some cases eyewear adversely affects performance (although many iris devices can scan people wearing glasses or contact lenses).
- Although the iris is a relatively stable biometric, it is affected by aging and disease.
- Relies on proprietary hardware and software technologies.
- Costs tend to be high compared to finger scanning, hand scanning, and many facial recognition systems.

On the other hand, the per unit cost of the leading hardware/software combination technology has dropped to as low as $300 per seat, still higher than finger scans but significantly lower than the over $5,000-per-seat price seen a few years ago.

29.4.5 Voice Recognition. Voice recognition biometrics "utilizes the distinctive aspects of the voice to verify the identity of individuals."[28] Voice recognition

generally is classified as a behavioral biometric, although it actually combines elements of behavioral and physiological biometrics: "The shape of the vocal tract determines to a large degree how a voice sounds, a user's behavior determines what is spoken and in what fashion."[29] Stated somewhat differently, "voice is a behavioral biometric but is dependent on underlying physical traits, which govern the type of speech signals we are able and likely to utter."[30] Because of comparatively low levels of accuracy and considerable user variability in voice dynamics, this biometric generally is used only for verification, not identification. Commonly deployed voice recognition systems can be divided into two types: text-dependent systems (the speaker is prompted to say a specific thing) and text-independent systems (the authentication system processes any utterances of the speaker), which provide a higher level of security because they are more difficult to spoof and provide better accuracy than text-dependent systems.

Strengths of voice recognition include:

- Capacity to leverage existing telephony infrastructure (as well as built-in computer microphones).
- Low cost when existing infrastructure is used.
- Ease of use.
- Interface with speech recognition and verbal passwords.
- High level of user acceptance. (This biometric does not suffer from the negative perceptions associated with all of the other leading biometrics.)

Weaknesses of voice recognition include:

- More susceptible to replay attacks than other biometrics.
- Accuracy levels are low compared to iris scanning, finger scans, and hand scans.
- Accuracy levels are negatively affected by ambient noise and low-quality capture devices.
- Accuracy, security, and reliability are challenged by individual variations in voice, such as speaking softly or loudly, hoarseness or nasality because of a cold, and so on.
- The stability of the biometric is affected by illness, aging, and other user behaviors including smoking.

29.4.6 Other Biometric Technologies. The five major biometric technologies just discussed collectively comprise the vast majority of biometric technology under deployment. Other biometric technologies that register on market share breakdowns are two of the behavioral type: signature scan and keystroke scan. Although both of these behavioral biometrics are well accepted (signature scanning more so than keystroke scanning), their usefulness is limited by their lack of accuracy.

Other behavioral biometrics under investigation include gait and lip motion. One physiological biometric that has received considerable attention because of its high accuracy and security rates is retinal scanning. However, most analysts believe that the problems associated with retinal scanning (lack of user acceptance, high cost, difficult and painful acquisition process) outweigh any advantages to this biometric. The consensus seems to be that iris scanning has replaced retinal scanning as the eye scanning biometric of choice.

The use of DNA as a biometric identifier has also been investigated, although it has significant weaknesses including the fact that DNA in body tissues (e.g., epithelial cells) can be obtained surreptitiously and transferred easily for nefarious purposes whereas the official methods of collection (e.g., taking blood samples) are relatively intrusive.

Other physiological biometrics that may prove useful in the future include body odor, skin reflectance, and ear shape.

29.5 TYPES OF ERRORS AND SYSTEM METRICS. All types of identification and authentication systems suffer from two types of errors: false accepts and false rejects.

29.5.1 False Accept. Also known as false match, false positive, or type 1 error, false accept is the likelihood, expressed as a percentage, that an imposter will be matched to a valid user's biometric. In some systems, such as those that attempt to secure entry to a weapons facility, a bank vault, or a high-level system administrator account, the false match/false accept rate is the most important metric to watch. In other systems, such as a facial recognition system deployed by a casino in an effort to spot card counters, a high level of false matches may be tolerated.

29.5.2 False Reject. Also known as false nonmatch, false negative, or type 2 error, false reject is the probability that "a user's template will be incorrectly judged to *not* match his or her enrollment template."[31] False nonmatches typically result in the user being locked out of the system. These false nonmatches can occur because of changes in a user's biometric data, changes in how the biometric data is presented, and/or changes in the environment. Biometric systems are generally more susceptible to false rejects than they are to false accepts.

29.5.3 Crossover Error Rate. An important metric in biometric systems is the *crossover error rate* (CER), also known as the equal error rate (EER). This useful metric is the intersection of the false accept and false reject rates. In general, a lower CER indicates the biometric device is more accurate and reliable than another biometric device with a higher CER. Exhibit 29.2 provides a summary of benchmark test–based accuracy/error rates for the five most prevalent biometric technologies. Each biometric technology is rank-ordered from most accurate to least accurate based on CER.

EXHIBIT 29.2 Accuracy/Error Rates of Leading Biometric Technologies

Biometric	False Match Rate	False No-Match Rate
Iris scan	0.0001%	1.1–1.4%
Finger scan	0.02%	.6%
Hand scan	0.3%	3.0%
Voice (text independent)	2–5%	5–10%
Voice (text dependent)	2.0%	0.03%
Face scan	.1%	1–2.5%

Source: Based on data contained in Anil K. Jain (2008), "Biometric Authentication," *Scholarpedia*, 3(6); 3716.

29.5.4 Failure to Enroll. Another critical metric in biometric systems is FTE. As Ashbourn explains, FTE refers to "a situation whereupon an individual is unable to enroll their biometric in order to create a template of suitable quality for subsequent automated operation."[32] Common reasons for failure to enroll include physical disability and a user whose physiological/behavioral characteristics are less distinctive than average. Nanavati, Thieme, and Nanavati observe that failure to enroll can be a major problem in "internal, employee-facing deployments" in which "high FTE rates are directly linked to increased security risks and increased system costs."[33] A final important metric is the "transaction time." Transaction time refers to "a theoretical time taken to match the live template against a reference sample."[34]

29.6 DISADVANTAGES AND PROBLEMS

29.6.1 General Considerations. Despite the many advantages over other commonly used authentication systems, the implementation of biometric authentication controls carries a number of risks and disadvantages. Even the most accurate biometric system is not perfect, and errors will occur. The error rates and the types of errors will vary with specific biometrics deployed and the circumstances of deployment. Certain types of errors, such as false matches, may pose fundamental, critical risks to organizational security. Other types of errors—failure to enroll, false nonmatch—may reduce organizational productivity and efficiency and increase costs. Organizations planning biometrics implementation will need to consider the acceptable error threshold. In any event, organizations deploying biometric authentication systems must not be lulled into a belief that they are invulnerable to errors and/or fraud. Certain biometric systems (e.g., iris scanning) are fairly impervious to fraud, while others (especially behavior-based systems) are much more susceptible to it. Facial scanning systems can be foiled with clothing, makeup, eyeglasses, and/or changes in hairstyle. Even relatively stable physiology-based biometrics like fingerprint scans can be defrauded with the use of rubber or gelatin fingers. Matsumoto outlines a gummy finger approach designed to fool even those countermeasures mentioned in Section 29.4.1.[35] The protein used has a similar galvanic response to flesh and, since it is very thin and attached to a live finger, has the correct temperature. In some cases, blowing warm air over the scanner may even raise the latent print of the intruder's predecessor.

The deployment of commonly used authentication systems (i.e., ID badges, passwords, etc.) requires relatively little training, although one could argue that better training on the development and use of passwords would improve security. This limited need for training is not the case with most of the most commonly used biometric systems. Both systems administrators and users need instruction and training to ensure smooth operation of the system. Some biometric systems are exquisitely sensitive to intra- and inter-user variation in presentation and performance. Their effectiveness becomes substantially compromised and error rates substantially increase in cases of significant variation and/or irregular presentation. A related problem concerns user acceptance of the biometric system. Some users may object to the deployment of biometrics due to concerns over privacy and intrusiveness. In other instances, users may object to the deployment of biometrics and avoid optimal interface with the system because of safety and/or health concerns, general fears, and/or cultural and religious beliefs. For example, some individuals may be concerned that biometric systems that require them to touch a finger pad or hand pad will unnecessarily expose them to germs and place them at risk for illness. Some users may fear that eye scans will damage their

eyes. Other users may object to eye scans on the basis that the eyes are the window to the soul. Anderson notes that some persons may object to the use of biometrics due to a personal interpretation of religious doctrine.[36] Notwithstanding users' beliefs and perceptions about the biometric system, in many cases features or elements related to the users and/or the operating environment will influence the successful implementation and effectiveness of the system.

29.6.2 Health and Disability Considerations. Individuals with arthritis and/or certain other disabilities and physical limitations may be unable to enroll in systems and subsequently, to align themselves physically in an optimal position with respect to biometric sensors. For example, users with severe hand arthritis may be unable to place their hand firmly as required on the hand geometry sensor, and users with migraines and associated photophobia may find it physically too uncomfortable to look straight into the light sensor for the iris scan. Some disabled people may have to be excluded from biometric systems altogether. Some relatively minor disabilities, such as a slight tremor, may compromise a legitimate user's ability to gain access through certain biometric systems. Variations in physical size can also influence system accuracy. An iris scanner positioned for a standard height range may fail to capture images of either very short or very tall individuals, or in some cases an individual's hands or fingers may be either too large or too small to be read accurately in a hand or finger scanner. Likewise, individuals with neck and back problems may find it difficult to use some biometric devices, depending on the kind of positioning required. Systems that rely on behavioral biometrics such as voice or signature are particularly vulnerable to variations and irregularities in user characteristics. For example, users who speak too softly, too loudly, or too rapidly may cause system errors. Minor changes in users' health can affect some biometric readings. Excessive skin moisture or lack of skin moisture can impact finger scans.

Although one of the ideal properties of a biometric is its universality, in reality not everyone has the characteristic or has it to the same degree. For example, some people are born without distinct fingerprints. In other cases, users may have lost the distinctiveness of their fingerprints because of years of manual labor, use of certain chemicals, scarring, or the aging process. Anderson notes that "people with dark-colored eyes and large pupils give poorer iris codes."[37] Certain eye diseases and metabolic conditions may also reduce or negate the efficacy of eye scan authentication. Age has a significant impact on the user-biometric-system interface. Definite physiological changes are associated with the aging process and can result in poor template matching with the live biometric. In this case, reenrollment may be needed. Fingerprints are affected by the aging process as the skin becomes drier and more brittle; voice patterns change in tonal quality over time; and facial shape or appearance may shift with age. Overall, the acceptability of a biometric system will be lessened if there is the impression that implementation of the system discriminates against, or has an otherwise adverse impact on, the disabled, the ill, ethnic minorities, the elderly, and other protected or traditionally disadvantaged groups of users. Organizations must ensure compliance with the Americans with Disabilities Act when implementing biometric authentication systems. Compliance may involve providing alternative methods of authentication to those affected.

29.6.3 Environmental and Cultural Considerations. A broad range of factors in the operating environment can also impact the effectiveness and acceptability of biometric systems. User-related cultural, social, and behavioral factors can influence

system performance. For instance, the accuracy of facial scans can be compromised by users' changes in hairstyle, facial hair, and headwear as well as by changes in an individual's physical appearance because of significant weight gain or loss. The accuracy of voice/speech recognition systems is affected by the distance between the scanner and the user as well as by the volume of speech. Fingerprint recognition is impeded in cases when users' skin is too dry, whether the condition arises as a result of aging, skin disease, environmental factors, or occupation-related factors, such as frequent hand washing among healthcare professionals. Factors in the surrounding ambient environment may also affect the accuracy of the biometric system. Ambient lighting will influence accuracy and error rate in facial scans and, to a lesser extent, in iris scans. Noise levels can impede the effectiveness of voice recognition systems. Humidity and air temperature can affect the accuracy of fingerprint and hand scans.

29.6.4 Cost Considerations. Although the cost of biometric system implementation has fallen dramatically in the past few years, it is still a major barrier for many organizations. Costs vary significantly depending on the type of system. Recent reports suggest that newer fingerprint scanners can be purchased for as little as $50 per unit; voice recognition systems can cost in excess of $50,000. However, even the least expensive biometrics systems are likely to cost more than simpler versions of traditional authentication systems. Experts estimate minimum costs, including hardware and software, at $200 or so per user and upward of $150,000 for corporate-wide protection in a medium-sized business. Compounding the cost issues are problems related to the lack of clear standards and the lack of clear interoperability between various biometric authentication systems.

Many of the problems and difficulties with biometrics systems are likely to be corrected or significantly mitigated with technological improvements, better user and administrator training, and good control of environmental conditions. In other cases, problems can be overcome or ameliorated with the use of countermeasures, such as combining different types of biometrics, combining biometrics with traditional authentication systems, and so on. Two major concerns that will continue to loom large and deserve closer examination are biometric identity theft and user privacy.

29.6.5 Attacks on Biometric Systems. Although biometrics are much less vulnerable to attack than other authentication controls, they are not immune to fraud. Moreover, when a biometric identity is stolen or spoofed, it creates a much bigger problem than that created by the theft of an ID badge, USB key, or password because a biometric cannot be simply canceled and replaced. One of the principal advantages to using biometrics for authentication is their invariability over time. Consequently, when an imposter or intruder defrauds a biometric authentication system and creates a false match error, the entire biometric security system is defrauded and the individual authorized user's biometric integrity is compromised. Likewise, Prabhakar, Pankanti, and Jain note, "One disadvantage of biometrics is that they cannot be easily revoked. If a biometric is ever compromised, it is compromised forever."[38]

A number of analysts believe that the ultimate solution to the problem of biometric identity theft lies in the development of "cancelable biometrics." Researchers at IBM have developed a prototype for the cancelable biometric that incorporates a repeatable distortion of the biometric. Similar in theory to the use of public and private keys for encryption, a unique distortion of the biometric is introduced at each enrollment. Therefore, if a user's biometric is compromised, only the one system is defrauded, not every system in which the user is enrolled.

29.6.6 Privacy Concerns. The use of biometric authentication controls raises significant privacy concerns, particularly in comparison to conventional authentication methods like passwords and ID badges. User objections to biometrics are often based on privacy concerns, sometimes articulated in terms of the user's sense of the intrusiveness of the biometric system. Anecdotal reports suggest that public perceptions of intrusiveness vary among different biometrics and in how biometrics are implemented. With regard to the latter, Nanavati, Thieme, and Nanavati[39] report that there is a greater risk of privacy invasiveness when:

- Deployment is covert (users are not aware of the system's operation) versus overt.
- The system is mandatory versus opt-in.
- The system is used for identification rather than verification.
- It is deployed for an indefinite duration versus fixed duration.
- It is deployed in the public versus the private sector.
- The user is interfacing with the system as an employee/citizen versus an individual/customer.
- An institution, not the user, owns the biometric information.
- The biometric data are stored in a template database versus the user's personal storage.
- The system stores identifiable biometric data versus templates.

A vivid example of the public's lack of acceptance of the covert use of biometric systems comes from the 2001 Super Bowl and the uproar that ensued after the Tampa Police Department deployed facial scanning technology for the purpose of picking out criminal suspects from the audience. In contrast, in the aftermath of the 2001 attacks on the World Trade Center and the Pentagon, there has been fairly widespread public acceptance of the use of facial scanning at airports in the United States.

Users generally view behavior-based biometrics, such as voice recognition and signature verification, as less intrusive and less privacy-threatening than physiology-based biometrics. Facial scanning is viewed as having a high potential for privacy invasion because of the capacity to deploy it without the user's knowledge and participation. Finger scans may be viewed as intrusive and privacy-invasive because of their association with law enforcement functions. The level of intrusiveness of the scanning technique appears to affect users' perception of privacy invasion, with iris scanning provoking more privacy objections than hand scanning. Civil libertarians and users also raise privacy objections over biometric systems that have the potential to uncover additional information about the user beyond the biometric identity. For example, finger scans, because of their capacity to be linked to large law enforcement databases of fingerprints, could be used to reveal information about the user's criminal background. Iris scans have the capacity to reveal confidential medical/health information about the user. Probably one of the most troubling privacy-related aspects of biometrics is the potential for large-scale linkage between biometric systems and the use of biometric data to facilitate large-scale national ID programs. Even though employers may design a biometric system for purely in-house use in order to facilitate verification of employee identities on corporate networks, federal regulations and laws such as the USA PATRIOT Act may eventually compel employers to surrender employees' private biometric data to government authorities.

In summary, the major privacy concerns associated with biometric deployments include:

- Users' loss of anonymity and autonomy
- Risk of unauthorized use of biometric information and/or unauthorized collection of biometric information
- Unnecessary collection of biometric information
- Unauthorized disclosure of biometric information to others
- Systematic reduction of users' reasonable expectation of privacy
- Potential for misuse on the part of overzealous or corrupt government agents

Many of these concerns can be generally lumped under the heading of "function creep" or "mission creep," wherein biometric systems designed for user authentication may, over time, be used for purposes not originally intended. An example of "mission creep" is the use of Social Security numbers (SSNs) for identification. The original Social Security cards were stamped "Not for Identification." However, many organizations (including the Internal Revenue Service) use SSNs for identification purposes.

Notwithstanding the privacy risks, supporters of biometric authentication systems argue that, properly deployed and with adequate best practice controls, biometric systems actually can function to enhance and protect privacy. Woodward, Orlans, and Higgins point out that "several newly developed biometric technologies use an individual's physical characteristics to construct a digital code for the individual without storing the actual physical characteristics," thus creating a sort of *biometric encryption* that can be used to protect the privacy of an individual's financial, medical, or other data.[40] Nanavati, Thieme, and Nanavati argue that "privacy-sympathetic" biometric systems can be designed.[41] Such systems would:

- Have limited system scope.
- Eschew use of biometrics as a unique identifier.
- Limit retention of biometric information.
- Limit storage of identifiable biometric data.
- Limit collection and storage of extraneous information, while including "opt-out" provisions for users.
- Enable anonymous enrollment and verification.
- Provide means of correcting and accessing biometric-related information.
- Limit system access.
- Use security tools and access policies to protect biometric information.
- Make provisions for third-party audits.
- Disclose the system purpose and objective.
- Disclose enrollment, verification, and identification processes.
- Disclose policies and protections in place to ensure privacy of biometric information.
- Disclose provisions for system termination.[42]

In contrast to this view, Alterman argues that the deployment of biometric systems and the use of biometric data for identification and verification are ethically questionable because they always entail a violation of privacy and autonomy. Alterman finds "something disturbing about the generalized use of biometric identification apart from the standard data privacy issue."[43] He maintains that biometric data "has inherent moral value"[44] but does not go so far as to argue against *any* deployment of biometric identification or verification systems. Rather, he maintains that they must be judiciously implemented and deployed only with due consideration to users' privacy concerns.

29.6.7 Legal Issues. Within the legal system, as mentioned, the question arises as to the validity of using biometrics. As many consider this to be an invasion of privacy on many levels, the actual definition of this should be briefly touched on.

According to current law, there are four separate categories of "invasion of privacy":

1. An intrusion into your private life—this is usually done by someone attempting to secretly learn something about you.
2. A public disclosure of private facts about yourself by someone other than yourself.
3. The use of your name, likeness, or both for monetary gain without your consent.
4. Any publication placing you in a false light.

With biometrics, the primary concern seems to stem from the first, and sometimes the second categories. Possibly these fears (and valid concerns, in light of recent government mandates and actions) may be alleviated by only storing a digital template, or a marker similar to those used in algorithms, rather than the actual image. TSA has been forced to use scans that show only a resemblance of the person in the scanner, and not an actual naked picture. The company they were using up to this time could not comply with their software, and thus lost the contract. With this, the possibility of the government being kept on a leash and not infringing on our Constitutional rights becomes a bit brighter.

29.7 RECENT TRENDS IN BIOMETRIC AUTHENTICATION

29.7.1 Government Advances in Biometric Authentication. Although private-sector organizations are increasingly adopting biometric technologies for their authentication needs, the government (public) sector has led investment in biometrics. The 2001 terrorist attacks on the World Trade Center and the Pentagon, and the ensuing USA PATRIOT Act, have encouraged increasing government commitment to biometric technologies. The Department of Defense (DoD), the Department of Homeland Security (DHS), the Immigration and Naturalization Service, and the Department of Transportation are the government agencies most involved in the deployment of biometrics technologies. The DoD's Common Access Card program involves putting biometric technology on a smart ID card.[45] The US-VISIT program under the DHS is another government program that incorporates biometrics (including face and fingerprint) into a smart ID card. Another DHS program, the Transportation Worker Identity Credential, incorporates biometric information in an ID card.

29.7.2 Face Scanning at Airports and Casinos. After the 2001 terrorist attacks on the World Trade Center and the Pentagon, most of the nation's airports

moved to incorporate face-scanning technologies into their security systems. Most studies of the effectiveness of these systems, however, have revealed their high error rates and low accuracy rates.

Casinos utilize facial scanning systems to identify professional "advantage players" and cheats. Although this is largely unregulated in the United States, Canadian casinos must notify players regarding the use of such systems. Casinos share data on professionals and cheats. One firm has networked 125 casino surveillance operators in the United States, Canada, Puerto Rico, Aruba, and the Bahamas and provides real-time alerts and other information useful in identifying suspicious players. However, it is unclear as to how such systems may be affected by international law. Article 12 of the United Nations Universal Declaration of Human Rights guarantees that "[n]o one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honor and reputation. Everyone has the right to the protection of the law against such interference or attacks."[46] Whether this system amounts to arbitrary interference or an attack on one's honor has not been addressed by the courts, but a case clearly could be made that covert use of such systems does so. A best practice for any surveillance system is informed consent. Organizations should clearly post notifications about the use of surveillance systems in order to protect themselves from legal challenges.

29.7.3 Increased Deployment in the Financial Industry. Usually slow to embrace new technologies, the financial industry is one of the leaders in the adoption of biometric authentication controls. Current deployments range from fingerprint scanners securing computer networks for brokers, to facial recognition systems at ATMs, to iris scanning for high-security access points. International Biometric Group projected that U.S. financial services firms would spend $672 million in 2007 for various biometric deployments. One of the biggest deployments to date has been United Bankers' Bancorporation (UBB) adoption of U.are.U, a fingerprint recognition system that allows UBB customers to automatically log onto UBB's Website with finger scans versus passwords. UBB also adopted a fingerprint authentication system for its employees. Wells Fargo, Bloomberg Financial, and Janus Capital Management are other well-known financial firms that have adopted biometric authentication systems for employees and/or customers. Although some financial institutions have selected voice, iris, or facial-scan-based systems, most seem to be choosing finger-scan systems.

29.7.4 Biometrics in the Healthcare Industry. Spurred in part by new regulations that require healthcare institutions to ensure the privacy and security of patient records, healthcare companies have also been at the forefront in the adoption of biometric authentication. Among the major healthcare organizations that have moved to biometric authentication is the Mayo Clinic, which adopted a fingerprint ID system in 2002. The majority of healthcare institutions that have adopted biometric authentication systems have selected finger-scan ID systems. However, deployment of these systems in healthcare organizations has not met with the same success as seen in the financial services industry. Issues involving the potential transmission of illness via physical contact with the fingerprint scanner are not trivial. Additionally, error rates have been higher and accuracy rates much lower than expected. The major reason behind the high incidence of errors appears to be the particulars of the healthcare environment, especially the characteristics of the hands of the doctors, nurses, and other healthcare workers using these systems. Specifically, system performance appears to be undermined by the chronically dry hands of these workers, a condition resulting from

frequent hand washing and the use of alcohol-based hand sanitizers. Another problem has been the resistance to using the fingerprint technology by both nurses and doctors, who feel that it involves a privacy intrusion. Add to this the costs that are associated with tying a single sign on type of authentication to the core hospital system, and the problems escalate even higher. Cerner and EPIC seem to have the top programs in the healthcare industry, and both have been certified for all three Meaningful Use levels required by HIPAA and HITECH. Yet, these programs come at a hefty cost. Even barebones, a recent change for a hospital in Wyoming cost them $4.5 million. While it was a change that needed to be done, it stretched their budgets to the very maximum, and several cuts were made to accommodate this.

29.7.5 Increased Deployment of Time and Attendance Systems. An increasing number of companies across many different industries are deploying biometric-based time-and-attendance systems. A shift from the past practice is in the increased use of biometric attendance and tracking systems for white-collar workers. Previously, the focus was on blue-collar factory workers. Although some employers are using the traditional hand-scanning systems, there appears to be a shift toward the use of finger-scanning time-and-attendance systems. This shift seems to be related to the more competitive pricing structure for the finger-scanning systems.

29.8 SUMMARY AND RECOMMENDATIONS.

There is no universal "best" biometric authentication system. Each of the five leading biometric technologies carries specific advantages and disadvantages. Some biometric technologies are more appropriate for certain applications and environments than their counterparts. An organization in the midst of evaluating potential biometrics authentication implementation must recognize that there will be trade-offs in any selection, such as cost for accuracy, privacy versus user acceptance, and so on, and there are not yet any universal decision factors for selecting a particular biometric technology for a specific application. There is, however, substantial research into many of the advantages and disadvantages of biometrics. Exhibit 29.3 provides a summary comparison of the features of the five leading biometric technologies discussed in this chapter. The features, shown in the extreme left column, were excerpted from various researcher efforts, and the rankings represent an amalgam of the rankings found in the literature.

Although biometric authentication systems promise cost savings and higher levels of security for organizations, they are not a panacea. Many factors affect how well or poorly biometric authentication controls will perform in a given organizational environment. Included among these factors are the users, the administration, the environment, the infrastructure, the budget, the communication system, and the existing security needs. Although many biometric technologies are capable of operating as stand-alone systems, in reality their accuracy and performance levels would be greatly improved by combining them with more conventional authentication methods, such as passwords and keys. Such multifactor systems offer greater security and reliability.

In selecting a biometric authentication system and preparing for an implementation, organizations should focus on the user-technology interface and the conditions in the operational environment that may influence the technology's performance. For example, the healthcare industry's unreflective embrace of finger scan technology illustrates the dangers of failing to heed environmental realities. It is important that organizations consider not only the practical impediments to effective implementation but also the potential psychological impediments, such as user fears about the technology. Ethically, the organization also has the obligation to consider carefully the extent to which the

EXHIBIT 29.3 Comparison of Leading Biometric Technologies

	Finger Scan	Facial Scan	Hand Scan	Iris Scan	Voice Recognition
Accuracy	High	Low	Medium	Very high	Low to medium
Ease of Use	High	Medium	High	Low to medium	High
User Acceptance	Medium	High (overt) Low (covert)	High	Low to medium	High
Privacy Concerns	High	Very high (overt)	Medium	High	Very low
Cost	Low to medium	Low to medium	Medium	High	Low
Performance	High	Low	Medium	High	Low
Potential for Circumvention	Medium	High	Low to medium	Very low	High
Distinctiveness	High	Low	Medium	Very high	Low
Barriers to Universality	Worn ridges; hand or finger impairment	None	Hand impairment	Visual impairment	Speech impairment
Susceptibility to Changes in Biometric	Low to medium	Medium to high	Medium	Low	Low to medium
Susceptibility to Changes in the Environment	Low	High	Very low	Low	Medium to high
Error-Causing Factors	Age, trauma, degradation of prints	Lighting, contrast, pose, movement, expression	Hand injury or trauma, inability to place correctly	Positioning, eye angle, glasses, disease	Illness, age, quality of communication system, ambient noise
Mitigations for Potential Errors	Periodic re-enrollment, enrollment of multiple fingers	Frequent re-enrollment, multiple scans, controlled environment	Periodic re-enrollment, enrollment of both hands	Periodic re-enrollment, user training, enroll both irises	Periodic re-enrollment, control ambient noise

implementation of biometric authentication compromises the privacy rights of users. In making this assessment, management must take into account the possibility that the organization may be compelled to release employees' biometric-related information to government authorities.

Recent studies in biometrics include banks and other financial institutions having identified the use of voice biometrics over fingerprint identification as one of the best means to secure its client accounts and financial information. Voice biometrics compares various characteristics drawn from a client's voice such as inflection, pitch, dialect, accent and others, matching it with data previously captured, and securely stored. For this technology to work, however, it will require banks and other financial institutions to register their clients' voice patterns, correlate them to personal data, and place this into a database.[47] Other studies have reached agreements with major corporations to help within the healthcare industry, such as BIO-Key having executed an OEM agreement with Caradigm for use of BIO-Key's Identity and Access Management suite. Caradigm will offer this suite within its Single Sign On solution to help hospitals and other caregivers writing electronic prescriptions of controlled substances the ability to meet federal and state regulations.[48] Other areas of study within this field have covered such areas as a wearable biometric tattoo for identification purposes, vein recognition, and mobile platforms being used to promote its acceptance. In a 2011 report released by the Unisys Corporation, acceptance of these and other biometric technologies are gaining, but more slowly than anticipated.[49] Future use and acceptance may come about as technology advances, people become more accustomed to their use in everyday life, and legislation forces the use for security purposes.

29.9 FURTHER READING

Kaine, A. K. "The Impact of Facial Recognition Systems on Business Practices within an Operational Setting." 25th International Conference Information Technology Interfaces, June 16–19, 2003, Cavtat, Croatia, pp. 315–320.

Hamilton, D. P. "Workplace Security; Read My Lips: Are Biometric Systems the Security Solution of the Future?" *Wall Street Journal*, September 29, 2003.

Jain, A. K, P. Flynn, and A. A. Ross, eds. *Handbook of Biometrics*. New York: Springer, 2007.

NIST Biometrics Resource Center Website, www.nist.gov/biometrics-portal.cfm

Ratha, N. K., Connell, J. H., & Bolle, R. M. "Enhanced Security and Privacy in Biometrics-Based Authentication Systems." *IBM Systems Journal* 40, No. 3 (2001): 614–634.

Ratha, N., and R. Bolle, eds. *Automatic Fingerprint Recognition Systems*. New York: Springer, 2003.

Ross, A. A., K. Nandakumar, and A. K. Jain. *Handbook of Multibiometrics*. New York: Springer, 2006.

Vacca, J. R. *Biometric Technologies and Verification Systems*. Boston: Butterworth-Heinemann, 2007.

29.10 NOTES

1. J. Boroshok, "Pointing the Finger at Biometrics." SearchSecurity Website, January 14, 2005, http://searchsecurity.techtarget.com/originalContent/0289142, sid14_gci1044805,00.html
2. International Biometric Group, "Biometrics Market and Industry Report 2004–2008," 2005, accessed June 4, 2005, www.biometricgroup.com/reports/public/market_report.html (URL inactive).

3. International Biometric Group, "Biometrics Market and Industry Report 2007–2012," January 2007, accessed February 4, 2007, www.biometricgroup.com/reports/public/market_report.html (URL inactive).
4. J. Feldman, "2013 IT Spending Priorities Survey," *Information Week*, May 6, 2013, http://reports.informationweek.com/abstract/83/10497/IT-Business-Strategy/Research:-2013-IT-Spending-Priorities-Survey.html
5. M. A. Davis, "2013 Strategic Security Survey," *Information Week*, May 27, 2013, www.informationweek.com/security/management/2013-strategic-security-survey/240155501
6. J. Chirillo and S. Blaul, *Implementing Biometric Security* (Indianapolis, IN: John Wiley & Sons, 2003), p. 2.
7. S. Nanavati, M. Thieme, and R. Nanavati, *Biometrics: Identity Verification in a Networked World* (Hoboken, NJ: John Wiley & Sons, 2002), p. 11.
8. A. K. Jain, "Biometric Recognition: How Do I Know Who You Are," *IEEE Symposia* (2004): 3–5.
9. R. M. Bolle, J. H. Connell, S. Pankanti, N. K. Ratha, and A. W. Senior, *Guide to Biometrics* (New York: Springer-Verlag, 2004).
10. Jain, "Biometric Recognition," p. 3.
11. Bolle et al., *Guide to Biometrics*, p. 6.
12. Jain, "Biometric Recognition," p. 3.
13. Chickowski, E. "Mobile Biometrics: Your Device Defines You", October 3, 2012, www.informationweek.com/quickview/mobile-biometrics-your-device-defines-yo/1675?queryText = biometrics&wc = 4
14. Nanavati et al., *Biometrics*, p. 12.
15. Nanavati et al., *Biometrics*, p. 14
16. Nanavati et al., *Biometrics*, p. 14.
17. Jain, "Biometric Recognition," p. 3.
18. J. D. Woodward, N. M. Orlans, and P. T. Higgins, *Biometrics* (New York: McGraw-Hill/Osborne, 2003), p. 37.
19. S. Liu and M. Silverman, "A Practical Guide to Biometric Security Technology," *IT Pro* (January/February 2001): 20.
20. Nanavati et al., *Biometrics*, p. 20.
21. Chirillo and Blaul, *Implementing Biometric Security*, p. 21.
22. Chirillo and Blaul, *Implementing Biometric Security*, p. 24.
23. Bolle et al., *Guide to Biometrics*, p. 36.
24. N. Imparato, "Does Face Recognition Have a Future?" *Intelligent Enterprise* 5, No. 7 (April 2002): 20.
25. H. Bray, "Facial Recognition Software Precision May Be Years Away," April 29, 2013, www.bostonglobe.com/business/2013/04/28/facial-recognition-technology-after-marathon-bombings/VN9gyIUqtps5EVfjI47P1 J/story.html
26. Nanavati et al., *Biometrics*, p. 99.
27. Brandom, R. "We Know Who You Are: The Scary New Technology of Iris Scanners," May 2, 2013, www.theverge.com/2013/5/2/4270352/theyre-already-watching-the-scary-new-technology-of-iris-scanners
28. Nanavati et al., *Biometrics*, p. 87.
29. Nanavati et al., *Biometrics*, p. 87.
30. Bolle et al., *Guide to Biometrics*, p. 40.

31. Nanavati et al., *Biometrics*, p. 27.
32. J. Ashbourn, P*ractical Biometrics: From Aspiration to Implementation* (London: Springer-Verlag, 2003), p. 10.
33. Nanavati et al., *Biometrics*, p. 35.
34. Ashbourn, *Practical Biometrics*, p. 10.
35. T. Matsumoto, "Importance of Open Discussion on Adversarial Analyses for Mobile Security Technologies: A Case Study for User Identification," May 14, 2002, http://crypto.csail.mit.edu/classes/6.857/papers/gummy-slides.pdf
36. R. J. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems* (Hoboken, NJ: John Wiley & Sons, 2001).
37. Anderson, *Security Engineering*, p. 274.
38. S. Prabhakar, S. Pankanti, and A. K. Jain, "Biometric Recognition: Security and Privacy Concerns," *IEEE Security & Privacy* (March/April 2003): 39.
39. Nanavati et al., *Biometrics*.
40. Woodward et al., *Biometrics*. p. 211.
41. Nanavati et al., *Biometrics*.
42. Nanavati et al., *Biometrics*.
43. A. Alterman, " 'A Piece of Yourself': Ethical Issues in Biometric Identification," *Ethics and Information Technology* 5, No. 3 (2003): 143.
44. Alterman, " 'A Piece of Yourself,' " p. 145.
45. www.digitalpersona.com/products/UPOS.php
46. United Nations, "Universal Declaration of Human Rights," December 10, 1948, www.un.org/en/documents/udhr/index.shtml
47. R. King, "Biometric Research Note: Mobile Devices To Drive Bank Adoption Of Voice Biometrics," BiometricUpdate.com, January 1, 2013, www.biometricupdate.com/201301/mobile-devices-to-drive-bank-adoption-of-voice-biometrics
48. A. Vrankuli, "Caradigm, BIO-key International Ink OEM Agreement," Biometric Update.com, May 13, 2013, www.biometricupdate.com/201305/caradigm-bio-key-international-ink-oem-agreement
49. J. Trader, "Public Opinion on Biometric Technology Points to Wider Acceptance," M2SYS Blog on Biometric Technology, November 8, 2011, http://blog.m2sys.com/comments-on-recent-biometric-news-stories/public-opinion-on-biometric-technology-points-to-wider-acceptance

CHAPTER 30

E-COMMERCE AND WEB SERVER SAFEGUARDS

Robert Gezelter

30.1 INTRODUCTION 30·2

30.2 BUSINESS POLICIES AND STRATEGIES 30·3

30.2.1 Step 1: Define Information Security Concerns Specific to the Application 30·3
30.2.2 Step 2: Develop Security Service Options 30·5
30.2.3 Step 3: Select Security Service Options Based on Requirements 30·7
30.2.4 Step 4: Ensures the Ongoing Attention to Changes in Technologies and Requirements 30·9
30.2.5 Using the Security Services Framework 30·9
30.2.6 Framework Conclusion 30·17

30.3 RULES OF ENGAGEMENT 30·18

30.3.1 Website-Specific Measures 30·18
30.3.2 Defining Attacks 30·19
30.3.3 Defining Protection 30·19
30.3.4 Maintaining Privacy 30·20
30.3.5 Working with Law Enforcement 30·20
30.3.6 Accepting Losses 30·20
30.3.7 Avoiding Overreaction 30·21
30.3.8 Appropriate Responses to Attacks 30·21
30.3.9 Counter-Battery 30·22
30.3.10 Hold Harmless 30·22

30.4 RISK ANALYSIS 30·22

30.4.1 Business Loss 30·23
30.4.2 PR Image 30·23
30.4.3 Loss of Customers/ Business 30·24
30.4.4 Interruptions 30·24
30.4.5 Proactive versus Reactive Threats 30·25
30.4.6 Threat and Hazard Assessment 30·25

30.5 OPERATIONAL REQUIREMENTS 30·25

30.5.1 Ubiquitous Internet Protocol Networking 30·26
30.5.2 Internal Partitions 30·27
30.5.3 Critical Availability 30·27
30.5.4 Accessibility 30·27
30.5.5 Applications Design 30·28
30.5.6 Provisioning 30·28
30.5.7 Restrictions 30·29
30.5.8 Multiple Security Domains 30·30
30.5.9 What Needs to Be Exposed? 30·30
30.5.10 Access Controls 30·31

30.5.11 Site Maintenance 30·32
30.5.12 Maintaining Site Integrity 30·32

30.6 TECHNICAL ISSUES 30·33

30.6.1 Inside/Outside 30·33
30.6.2 Hidden Subnets 30·34
30.6.3 What Need Be Exposed? 30·34
30.6.4 Multiple Security Domains 30·35
30.6.5 Compartmentalization 30·37
30.6.6 Need to Access 30·38
30.6.7 Accountability 30·39
30.6.8 Read-Only File Security 30·39
30.6.9 Going Offline 30·40
30.6.10 Auditing 30·40
30.6.11 Emerging Technologies 30·41

30.7 ETHICAL AND LEGAL ISSUES 30·41

30.7.1 Liabilities 30·41
30.7.2 Customer Monitoring, Privacy, and Disclosure 30·42
30.7.3 Litigation 30·43
30.7.4 Application Service Providers 30·44

30.8 SUMMARY 30·45

30.9 FURTHER READING 30·46

30.10 NOTES 30·46

30.1 INTRODUCTION. Today, electronic commerce involves the entire enterprise. The most obvious e-commerce applications involve business transactions with outside customers on mobile devices and the World Wide Web (WWW or Web), yet they are merely the proverbial tip of the iceberg. The presence of e-commerce has become far more pervasive, often involving the entire logistical and financial supply chains that are the foundations of modern commerce. Even the smallest organizations now rely on the Web for access to services and information.

The explosive growth of mobile applications in the past few years has integrated electronic transactions deeply in the minute-to-minute fabric of daily life. Walk down the street in any urban area, mobile devices have become ubiquitous. Nations which lacked wired infrastructure have skipped a technological generation and gone directly to a telecommunications infrastructure based on mobile phone technologies.

The pervasive desire to improve efficiency often causes a convergence between the systems supporting conventional operations with those supporting an organization's online business. It is thus common for internal systems at bricks-and-mortar stores to utilize the same back-office systems as are used by Web customers. It is also common for kiosks and cash registers to use wireless local area networks to establish connections back to internal systems. These interconnections can provide intruders with access directly into the heart of the enterprise.

The TJX case, which came to public attention in the beginning of 2007, was one of a series of large-scale compromises of electronically stored information on back-office and e-commerce systems. Most notably, the TJX case appears to have started with an insufficiently secured corporate network and the associated back-office systems, not a Website penetration. This breach escalated into a security breach of corporate data systems. It has been reported that at least 94 million credit cards were compromised.[1] On November 30, 2007, it was reported that TJX, the parent organization of stores including TJ Maxx and Marshall's, agreed to settle bank claims related to VISA cards for US$ 40.9M.[2]

E-commerce has now come of age, giving rise to fiduciary risks that are important to senior management and to the board of directors. The security of data networks, both those used by customers and those used internally, now has reached the level where it significantly affects the bottom line. TJX has suffered both monetarily and in public relations, with stories concerning the details of this case appearing in the *Wall Street*

Journal, the *New York Times*, *Business Week*, and many industry trade publications. Data security is no longer an abstract issue of concern only to technology personnel. The legal settlements are far in excess of the costs directly associated with curing the technical problem.

Protecting e-commerce information requires a multifaceted approach, involving business policies and strategies as well as the technical issues more familiar to information security professionals.

Throughout the enterprise, people and information are physically safeguarded. Even the smallest organizations have a locked door and a receptionist to keep outsiders from entering the premises. The larger the organization, the more elaborate the precautions needed. Small businesses have simple locked doors; larger enterprises often have many levels of security, including electronic locks, security guards, and additional levels of receptionists. Companies also jealously guard the privacy of their executive conversations and research projects. Despite these norms, it is not unusual to find that information security practices are weaker than physical security measures. Connection to the Internet (and within the company, to the intranet) worsens the problem by greatly increasing the risk and decreasing the difficulty, of attacks. Ubiquitous mobile devices only serve to compound the difficulties.

30.2 BUSINESS POLICIES AND STRATEGIES. In the complex world of e-commerce security, best practices are constantly evolving. New protocols and products are announced regularly. Before the Internet explosion, most companies rarely shared their data and their propriety applications with any external entities, and information security was not a high priority. Now companies taking advantage of e-commerce need sound security architectures for virtually all applications. Effective information security has become a major business issue. This chapter provides a flexible framework for building secure e-commerce applications and assistance in identifying the appropriate and required security services. The theoretical examples shown are included to facilitate the reader's understanding of the framework in a business-to-customer (B2C) and business-to-business (B2B) environment.

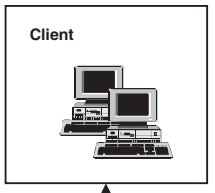
A reasonable framework for e-commerce security is one that:

1. Defines information security concerns specific to the application.
2. Defines the security services needed to address the security concerns.
3. Selects security services based on a cost-benefit analysis and risk versus reward issues.
4. Ensures the ongoing attention to changes in technologies and requirements as both threats and application requirements change.

This four-step approach is recommended to define the security services selection and decision-making processes.

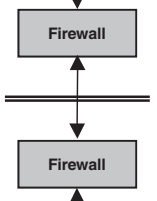
30.2.1 Step 1: Define Information Security Concerns Specific to the Application. The first step is to define or develop the application architecture and the data classification involved in each transaction. This step considers how the application will function. As a general rule, if security issues are defined in terms of the impact on the business, it will be easier to discuss with management and easier to define security requirements.

The recommended approach is to develop a transactional follow-the-flow diagram that tracks transactions and data types through the various servers and networks. This should be a functional and logical view of how the application is going to work—that is, how transactions will occur, what systems will participate in the transaction management, and where these systems will support the business objectives and the organization's product value chain. Data sources and data interfaces need to be identified, and the information processed needs to be classified. In this way a complete transactional flow can be represented. (See Exhibit 30.1.)



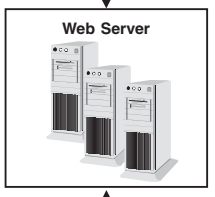
Presentation Layer

Trusted Communication	Authentication/ Identification	Audit	Access Control
• HTTP • HTTPS • Certificates • SSL Encryption • WAP • TLS	• None Required	• None Required	• None Required



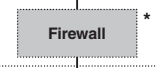
Network Layer

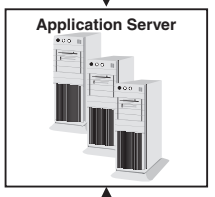
Trusted Communication	Authentication/ Identification	Audit	Access Control
VPN	None Required	None Required	None Required



Middle Layer

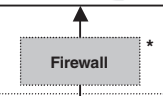
Trusted Communication	Authentication/ Identification	Audit	Access Control
• HTTP • HTTPS • SSL • SET • IOTP	• Application Password • PKI	• Transaction Logs • OS Logs	• OS • Application (RBAM)

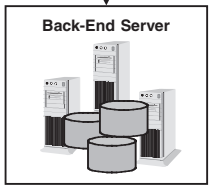




Application Layer

Trusted Communication	Authentication/ Identification	Audit	Access Control
• TCP/IP • IPSec	• App • OS • DB	• OS • APP • Database	• OS • DB Database (App Password)





Internal Layer

Trusted Communication	Authentication/ Identification	Audit	Access Control
• TCP/IP • IPSec	• OS • Database	• OS • Database	• OS • Database

*Firewall should be between the web layer/application layer or the application/back-end layers depending on the company's architecture.

EXHIBIT 30.1 Trust Levels for B2C Security Services

Common tiered architecture points include:

- **Clients.** These may be PCs, thin clients (devices that use shared applications from a server and have small amounts of memory), personal digital assistants (PDAs), and wireless application protocol (WAP) telephones.
- **Servers.** These may include World Wide Web, application, database, and middleware processors, as well as back-end servers and legacy systems.
- **Network devices.** Switches, routers firewalls, NICs, codecs, modems, and internal and external hosting sites.
- **Network spaces.** Network demilitarized zones (DMZs), intranets, extranets, and the Internet.

It is important at this step of the process to identify the criticality of the application to the business and the overriding security concerns: transactional confidentiality, transactional integrity, or transactional availability. Defining these security issues will help justify the security services selected to protect the system. The more completely the architecture can be described, the more thoroughly the information can be protected via security services.

30.2.2 Step 2: Develop Security Service Options. The second step considers the security services alternatives for each architecture component and the data involved in each transaction. Each architectural component and data point should be analyzed and possible security services defined for each. Cost and feasibility should not be considered to any great degree at this stage. The objective is to form a complete list of security service options with all alternatives considered. The process should be comparable with, or use the same techniques as, brainstorming. All ideas, even if impractical or farfetched, should be included.

Decisions should not be made during this step; that process is reserved for Step 3.

The information security organization provides services to an enterprise. The services provided by information security organizations vary from company to company. Several factors will determine the required services, but the most significant considerations include:

- Industry factors
- The company's risk appetite
- Maturity of the security function
- Organizational approach (centralized or decentralized)
- Impact of past security incidents
- Internal organizational factors
- Political factors
- Regulatory factors
- Perceived strategic value of information security

Several factors contribute to the services that information security organizations provide. "Security services" are defined as safeguards and control measures to protect the confidentiality, integrity, and accountability of information and computing resources. Security services that are required to secure e-commerce transactions need to be based on the business requirements and on the willingness to assume or reduce the risk of the

information being compromised. Information security professionals can be subject-matter experts, but they are rarely equipped to make the business decisions required to select the necessary services. Twelve security services that are critical for successful e-commerce security have been identified:

1. **Policy and procedures** are a security service that defines the amount of information security that the organization requires and how it will be implemented. Effective policy and procedures will dovetail with system strategy, development, implementation, and operation. Each organization will have different policies and procedures; best practice dictates that organizations have policies and procedures based on the risk the organization is willing to take with its information. At a minimum, organizations should have a high-level policy that dictates the proper use of information assets and the ramifications of misuse.
2. **Confidentiality and encryption** are a security service that secures data while they are stored or in transit from one machine to another. A number of encryption schemes and products exist; each organization needs to identify those products that best integrate with the application being deployed. For a discussion of cryptography, see Chapters 7 and 37 in this *Handbook*.
3. **Authentication and identification** are a security service that differentiates users and verifies that they are who they claim to be. Typically, passwords are used, but stronger methods include tokens, smart cards, and biometrics. These stronger methods verify what you have (e.g., token) or who you are (e.g., biometrics), not just what you know (password). Two-factor authentication combines two of these three methods and is referred to as strong authentication. For more on this subject, see Chapter 28 in this *Handbook.*
4. **Authorization** determines what access privileges a user requires within the system. Access includes data, operating system, transactional functions, and processes. Access should be approved by management who own or understand the system before access is granted. Authorized users should be able to access only the information they require for their jobs.
5. **Authenticity** is a security service that validates a transaction and binds the transaction to a single accountable person or entity. Also called nonrepudiation, authenticity ensures that a person cannot dispute the details of a transaction. This is especially useful for contract and legal purposes.
6. **Monitoring and audit** provide an electronic trail for a historical record of the transaction. Audit logs consist of operating system logs, application transaction logs, database logs, and network traffic logs. Monitoring these logs for unauthorized events is considered a best practice.
7. **Access controls and intrusion detection** are technical, physical, and administrative services that prevent unauthorized access to hardware, software, or information. Data are protected from alteration, theft, or destruction. Access controls are preventive—stopping unauthorized access from occurring. Intrusion detection catches unauthorized access after it has occurred, so that damage can be minimized and access cut off. These controls are especially necessary when confidential or critical information is being processed.
8. **Trusted communication** is a security service that assures that communication is secure. In most instances involving the Internet, this means that the communication will be encrypted. In the past, communication was trusted because it

was contained within an organization's perimeter. Communication is currently ubiquitous and can come from almost anywhere, including extranets and the Internet.

9. **Antivirus** is a security service that prevents, detects, and cleans viruses, Trojan horse programs, and other malware.
10. **System integrity controls** are security services that help to assure that the system has not been altered or tampered with by unauthorized access.
11. **Data retention and disposal** are a security service that keeps required information archived, or deletes data when they are no longer required. Availability of retained data is critical when an emergency exists. This is true whether the problem is a systems outage or a legal process, whether caused by a natural disaster or by a terrorist attack (e.g., September 11, 2001).
12. **Data classification** is a security service that identifies the sensitivity and confidentiality of information. The service provides guides for information labeling, and for protection during the information's life.

Once an e-commerce application has been identified, the team must identify the security issues with that specific application and the necessary security services. Not all of the services will be relevant, but using a complete list and excluding those that are not required will assure a comprehensive assessment of requirements, with appropriate security built into the system's development. In fact, management can reconcile the services accepted with their level of risk acceptance.

30.2.3 Step 3: Select Security Service Options Based on Requirements. The third step uses classical cost-benefit and risk-management analysis techniques to make a final selection of security service options. However, we recommend that all options identified in Step 3 be distributed along a continuum, such as shown in Exhibit 30.2, so that they can be viewed together, and compared.

Gauging and comparing the level of security for each security service and the data within the transaction will facilitate the decision process. Feasible alternatives can then be identified and the best solution selected based on the requirements. The most significant element to consider is the relative reduction in risk of each option, compared with the other alternatives. The cost-benefit analysis is based on the risk versus reward issues. The effectiveness information is very useful in a cost-benefit model.

Four additional concepts drive the security service option selection:

1. Implementation risk or feasibility
2. Cost to implement and support
3. Effectiveness in increasing control, thereby reducing risk
4. Data classification

Implementation risk considers the feasibility of implementing the security service option. Some security systems are difficult to implement due to factors such as product maturity, scalability, complexity, and supportability. Other factors to consider include skills available, legal issues, integration required, capabilities, prior experience, and limitations of the technology.

Cost to implement and support measures the costs of hardware and software implementation, support, and administration. Consideration of administration issues is

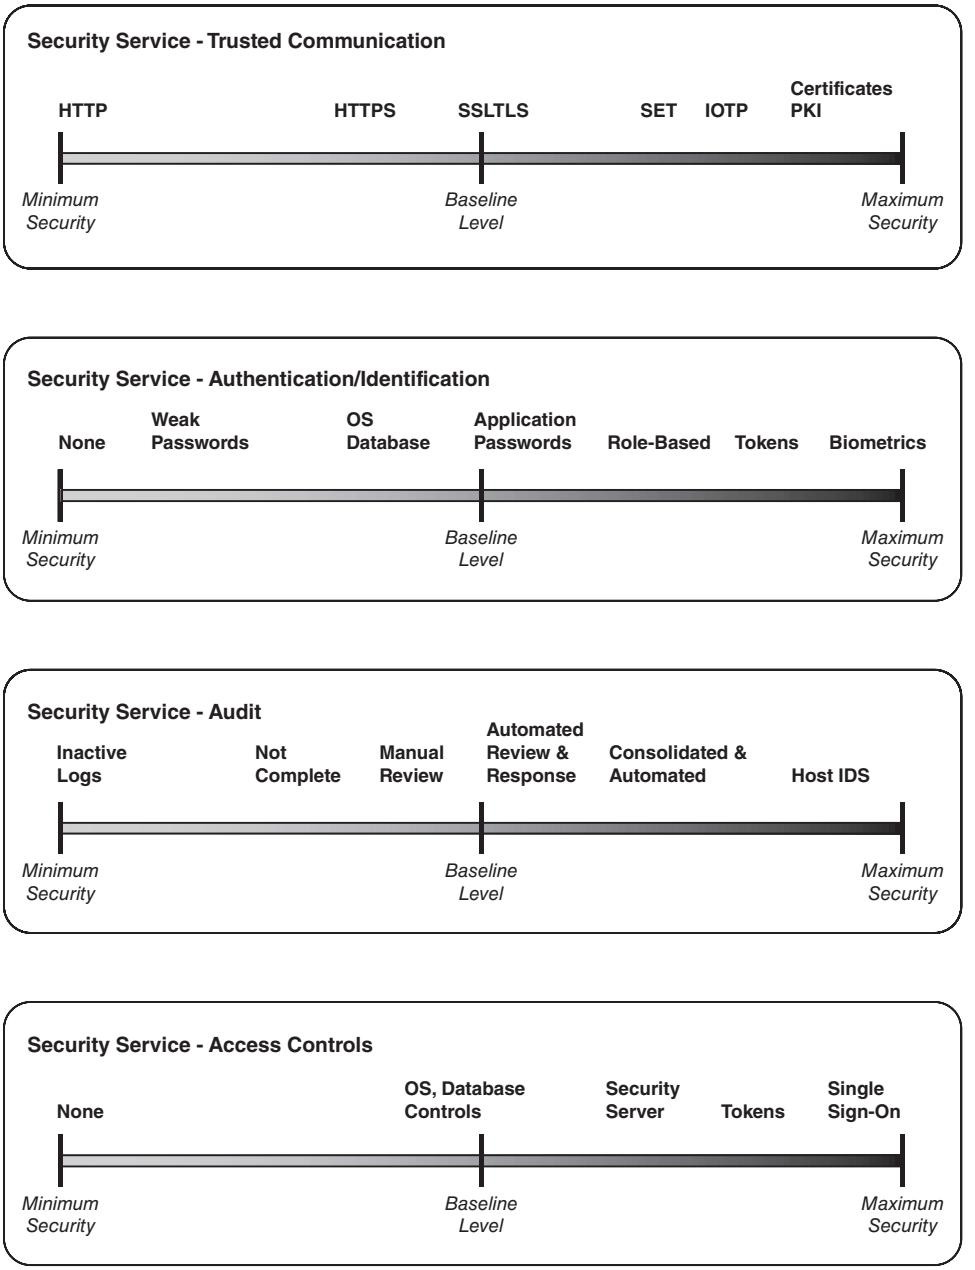


EXHIBIT 30.2 Continuum of Options

especially critical because high-level support of the security service is vital to an organization's success.

Effectiveness measures the reduction of risk proposed by a security service option once it is in production. Risk can be defined as the impact and likelihood of a negative event occurring after mitigating strategies have been implemented. An example of a negative event is the theft of credit card numbers from a business's database. Such an event causes not only possible losses to consumers but also negative public relations

that may impact future business. Effective security service options reduce the risk of a negative event occurring.

Data classification measures the sensitivity and confidentiality of the information being processed. Data must be classified and protected from misuse, disclosure, theft, or destruction, regardless of storage format, throughout their life (from creation through destruction). Usually, the originator of the information is considered to be the owner of the information and is responsible for classification, identification, and labeling. The more sensitive and confidential the information, the more information security measures will be required to safeguard and protect it.

30.2.4 Step 4: Ensures the Ongoing Attention to Changes in Technologies and Requirements. The only constant in this analysis is the need to evolve and to address ever-increasing threats and technologies. Whatever security approaches the preceding steps identify, they must always be considered in the context of the continuing need to update the selected approaches. Changes will be inevitable, whether they arrive from compliance, regulation, technological advances, or new threats.

30.2.5 Using the Security Services Framework. The next two sections are examples to demonstrate the power of the security services methodology. The first example is a B2C model; the business could be any direct-selling application. The second example is a B2B model. Both businesses take advantage of the Internet to improve their product value chain. The examples are a short demonstration of the security services methodology, neither complete nor representative of any particular application or business.

30.2.5.1 Business-to-Customer Security Services. The B2C company desires to contact customers directly through the Internet, and allow them to enter their information into the application. These assumptions are made to prepare this B2C system example:

- Internet-facing business
- Major transactions supported
- External customer-based system, specifically excluding support, administration, and operations
- Business-critical application
- Highly sensitive data classification
- Three-tiered architecture
- Untrusted clients, because anyone on the Internet can be a customer

Five layers must be secured:

1. The *presentation layer* is the customer interface, what the client sees or hears using the Web device. The client is the customer's untrusted PC or other device. The security requirements at this level are minimal because the company will normally not dictate the security of the customer. The identification and authentication done at the presentation level are those controls associated with access to the device. The proliferation of appliances in the client space (e.g., traditional

PCs, thin desktops, and PDAs) makes it difficult to establish a uniform access control procedure at this level. Although this layer is the terminal endpoint of the secure connection, it is not inherently trustworthy, as has been illustrated by all too many incidents involving public computers in cafes, hotels, and other establishments.

2. The *network layer* is the communication connection between the business and the customer. The client, or customer, uses an Internet connection to access the B2C Web applications. The security requirements are minimal, but sensitive, and confidential traffic will need to be encrypted.
3. The *middle layer* is the Web server that connects to the client's browser and can forward and receive information. The Web server supports the application by being an intermediary between the business and the customer. The Web server needs to be very secure. Theft, tampering, and fraudulent use of information needs to be prevented. Denial of service and Website defacement are also common risks that need to be prevented in the middle layer.
4. The *application layer* is where the information is processed. The application serves as an intermediary between the customer requests and the fulfillment systems internal to the business. In some examples, the application server and database server are the same because both the application and database reside on the same server. However, they could reside within the Web server in other cases.
5. The *internal layer* is comprised of the business's legacy systems and databases that support customer servicing. Back-end servers house the supporting application, including order processing, accounts receivable, inventory, distribution, and other systems.

For each of these five levels, we need four security services:

1. Trusted communications
2. Authentication/identification
3. Audit
4. Access controls

Step 1: Define Information Security Concerns Specific to the Application. Defining security issues will be particular to the system being implemented. To understand the risk of the system, the best starting place is with the business risk; then defining risks at each element of the architecture.

Business Risk

- The application needs high availability, because customers will demand contact at off-hours and on weekends.
- The Web pages need to be secure from tampering and cyber-vandalism, because the business is concerned about the loss of customer confidence as a result of negative publicity.
- Customers must be satisfied with the level of service.
- The system will process customer credit card information and will be subject to the privacy regulations of the Gramm-Leach-Bliley Act, 15 USC §§ 6801–09.

Technology Concerns
Of the five architectural layers in this example, four will need to be secured:

1. The *presentation layer* will not be secured or trusted. Communications between the client and the Web server will be encrypted at the *network layer*.
2. The *network layer* will need to filter unwanted traffic, to prevent denial of service (DoS) attacks and to monitor for possible intrusions.
3. The *middle layer* will need to prevent unauthorized access, be tamper-proof, contain effective monitoring, and support efficient and effective processing.
4. The *application layer* will need to prevent unauthorized access, support timely processing of transactions, provide effective audit trails, and process confidential information.
5. The *internal layer* will need to prevent unauthorized access, especially through Internet connections, and to protect confidential information during transmission, processing, and storage.

Step 2: Develop Security Services Options. The four security services reviewed in this example are the most critical in an e-commerce environment. Other services such as nonrepudiation and data classification are important but not included in order to simplify the example. Services elected are:

- Trusted communication
- Authentication and identification
- Monitoring and auditing
- Access control

Many security services options are available for the B2C case, with more products and protocols on the horizon. There are five architectural layers for each of the services defined in Step 1.

1. **Presentation layer.** Several different options can be selected for trusted communication. While the most common application-level protocol is the Hypertext Transfer Protocol (HTTP); the ever increasing need for privacy and assurance of integrity argues for the use of its more secure encrypted sibling, Transport Layer Security,[3] the successor to Secure Socket Layer (SSL) protocol,[4] in conjunction with certificates in a Public Key Infrastructure (PKI), to authenticate and cryptographically secure the communications pathway. The use, or possible use, of wireless connections somewhere in the path only increases the importance of secured, encrypted paths. Because the client is untrusted, the client's authentication, audit, or access control methods cannot be relied on.
2. **Network layer.** Virtual private networks (VPNs) are considered best practice for secure network layer communication. Firewalls are effective devices for securing network communication. The client may have a personal firewall. If the client is in a large corporation, there is a significant likelihood that a firewall will intervene in communications. If the client is using a home or laptop computer, then a personal firewall may protect traffic. There will also be a firewall on the B2C company side of the Internet.

3. **Middle layer.** The Web server and the application server security requirements are significant. Unauthorized access to the Web server can result in Website defacement by hackers who change the Web data. More important, access to the Web or application server can lead to theft, manipulation, or deletion of customer or proprietary data. Communication between the Web server and the client needs to be secure in e-commerce transactions. HTTP is the most common form of browser communication. In 2000, it was reported that over 33 percent of credit card transactions were using unsecured HTTP.[5] In 2007, VISA reported that compliance with the Payment Card Industry Data Security Standard (PCI DSS) had improved to 77 percent of the largest merchants in the United States,[6] still far from universal. This lack of encryption is a violation of both the merchant's agreement and the PCI, but episodes continue to occur. In the same vein, it is not unusual for organizations to misuse encryption, whether it involves self-signed, expired, or not generally recognized certificates. (See Chapter 37 in this *Handbook*.) HTTPS (HTTP over TLS) is the most common secure protocol solution, but the encryption key length and contents are critical: The larger the key, the more secure the transaction is from brute-force decryption (see Chapter 7). Digital certificates in a PKI and digital signatures are not as common, but they are more effective forms of security.
4. **Application layer.** The application server provides the main processing for the Web server. This layer may include transaction processing and database applications; it needs to be secure to prevent erroneous or fraudulent processing. Depending on the sensitivity of the information processed, the data may need to be encrypted. Interception and the unauthorized manipulation of data are the greatest risks in the application layer.
5. **Internal layer.** In a typical example, the internal layer is secured by a firewall protecting the external-facing system. This firewall helps the B2C company protect itself from Internet intrusions. Database and operating system passwords and access control measures are required.

Management can decide which security capability is required and at what level. This format can be repeated to discuss security services at all levels and all systems, not just e-commerce–related systems.

Step 3: Select Security Service Options Based on Requirements. In Step 3 the B2C company can analyze and select the security services that best meet its legal, business, and information security requirements. There are four stages required for this analysis:

1. Implementation risk or feasibility
2. Cost to implement and support
3. Effectiveness in increasing control, thereby reducing risk
4. Data classification

Implementation risk is a function of the organization's ability to effectively roll out the technology. In this example, we assume that implementation risk is low and resources are readily available to implement the technology.

Costs to implement and support are paramount to the decision-making process. Both costs need to be considered together. Unfortunately, the cost to support is difficult

to quantify and easily overlooked. In this example, resources are available to both implement and support the technology.

Effectiveness in increasing control is an integral part of the benefit and risk-management decisions. Each component needs to be considered in order to determine cross-benefits where controls overlap and supplement other controls. In this case, the control increase is understood and supported by management.

Data classification is the foundation for requirements. It will help drive the cost-benefit discussions because it captures the value of the information to the underlying business. In this example, the data are considered significant enough to warrant additional security measures to safeguard the data against misuse, theft, and loss.

There are many technology decisions required to secure the example environment. Management can use this approach to plot the security levels required by the system. For example, for system audit services, in order of minimal service to maximum:

- The minimal level of security is to have systems with a limited number of system events logged. For example, the default level of logging from the manufacturer is used but does not contain all of the required information. The logs are not reviewed but are available for forensics in the event of a security incident.
- A higher level of security is afforded with a log that records more activities based on requirements and not, for example, the manufacturer's default level. As in the minimal level of security, the activities are logged and available to support forensics but are not reviewed. In this case, more types of information are recorded in the system log, but it still may not contain all that is required.
- A sufficient log is kept on each server and is manually monitored for anomalies and potential security events.
- The log is automatically reviewed by software on each server.
- System logs are consolidated onto a centralized security server. Data from the system logs are transmitted to the centralized security server, and software is then used to scan the logs for specific events that require attention. Events such as attempts to gain escalated privileges to root or administrative access can be flagged for manual review.
- The maximum service level is a host-based intrusion detection system (IDS) used to scan the system logs for anomalies and possible security events. Once detected, action needs to be taken to resolve the intrusion. The procedure should include processes such as notification, escalation, and automated defensive response.

30.2.5.2 Business-to-Business Security Services. The second case study uses the security services framework in a B2B example. Following is a theoretical discussion of how the framework can be applied to B2B e-commerce security. These assumptions may be made in this B2B system example:

- Internet-facing.
- Supports major transactions.
- Descriptions will be external and customer based (excluding support, administration, and operations security services).
- Trusted communication is required.

- Three-tier architecture.
- Untrusted client.
- Business-critical application.
- Data are classified as highly sensitive.

There are five layers in this example that need to be secured:

1. The *presentation layer* is the customer interface and is what the client sees or hears using the Web device. The client is the customer's untrusted PC, but more security constraints can be applied because the business can dictate enhanced security. As noted previously, the security of the presentation layer is complicated by the wide range of potential client devices that may be employed.
2. The *application layer* is where the information is processed. The application serves as an intermediary between the business customer's requests and the fulfillment systems internal to the business (the back-end server). The application server is the supporting server and database.
3. The *customer internal layer* is the interface between the application server supporting the system at the customer's business location, and the customer's own internal legacy applications and systems.
4. The *network layer* is the communication connection between the business and another business. The Internet is used to connect the two businesses. Sensitive and confidential traffic will need to be encrypted. Best practice is to have the traffic further secured using a firewall.
5. The *internal layer* is the business's legacy systems that support customer servicing. The back-end server houses the supporting systems, including order processing, accounts receivable, inventory, distribution, and other systems.

The four security services are:

1. Trusted communications
2. Authentication/identification
3. Audit
4. Access controls

Step 1: Define Information Security Concerns Specific to the Application. Defining security issues will be particular to the system being implemented. To understand the risk of the system, it is best to start with the business risk; then define risk at each element of the architecture.

Business Risk

- Communication between application servers needs to be very secure. Data must not be tampered with, stolen, or misrouted.
- Availability is critical during normal business hours.
- Cost savings realized by switching from electronic data interchange (EDI) to the Internet is substantial, and will more than cover the costs of the system.

Technology Concerns
There are six architectural layers in this example, five of which need to be secured:

1. The *presentation layer* will not be secured or trusted. The communication between the client and the customer application is trusted because it uses the customer's private network.
2. The *application server* will need to be secure. Traffic between the two application servers will need to be encrypted. The application server is inside the customer's network and demonstrates a potentially high, and perhaps unnecessary degree of trust between the two companies (see Section 30.6.5).
3. The *customer's internal layer* will be secured by the customer.
4. The *network layer* needs to filter out traffic that is not required, prevent DoS attacks, and monitor for possible intrusions. Two firewalls are shown: one to protect the client and the other to protect the B2B company.
5. The *application layer* will need to prevent unauthorized access, support timely processing of transactions, provide effective audit trails, and process confidential information.
6. The *internal layer* will need to prevent unauthorized access (especially through Internet connections) and protect confidential information during transmission, processing, and storage.

Step 2: Develop Security Services Options. There are four security services reviewed in this example. Others could have been included, such as authenticity, nonrepudiation, and confidentiality, but they have been excluded to simplify this example. Elected security services include:

1. Trusted communication
2. Authentication/identification
3. Audit
4. Access control

Many security services options are available for B2B environments.

- **Presentation layer.** Several different options can be selected for communication. HTTP is the most common. The communications between the presentation layer residing on the client device and the application server, in this example, are internal to the customer's trusted internal network and will be secured by the customer. Past episodes involving this scenario include the previously mentioned TJX affair and the episode involving Hannaford.[7] Recently, Bashas', a chain of grocery stores in Arizona, was compromised by malware on their Point of Sale network.[8] These compromises of internal networks argue for the use of HTTPS (HTTP over TLS) as a presentation layer, even on "internal" networks where information transmitted may include personally identifiable information (PII) in some form. As a matter of safety, when in doubt it is far safer to run an encrypted infrastructure as a guard against exposing PII than it is to discover at a later date that some data in certain circumstances is protected PII. Encrypting unsensitive data does accrue excessive costs; failing to encrypt even a small amount of PII could lead to a significant liability.

- **Application layer.** Communication between the two application servers needs to be secure. The easiest and most secure method of peer-to-peer communication is via a VPN.
- **Customer internal.** Communications between the customer's application server and the customer's back-end server are internal to the customer's trusted internal network and will be secured by the customer.
- **Network layer.** It is common in a B2B environment that a trusted network is created via a VPN. The firewalls will probably participate in these communications, but hardware solutions are also possible.
- **Application layer.** The application server is at both the customer and B2B company sites. VPN is the most secure communication method. The application server also needs to communicate with the internal layer, and this traffic should be encrypted as well.
- **Internal layer.** The internal layer may be secured with another firewall from the external-facing system. This firewall helps the B2B company to protect itself from intrusions and unauthorized access. In this example, a firewall is not assumed, so the external firewall and DMZ need to be very secure.

Intrusion detection, log reading, and other devices can easily be added and discussed with management. This format can be repeated to discuss security services at all levels and all systems, not just e-commerce–related systems.

Step 3: Develop Security Service Options. In Step 3, the B2B company can analyze and select the security services that best meet its legal, business, and information security requirements. The biggest difference between B2C and B2B systems is that the B2C system assumes no level of trust. The B2B system assumes trust to some degree, but additional coordination and interface with the B2B customer or partner is required. Trustworthiness is not an either/or question and should not be presumed (see Section 30.6.4). This coordination and interoperability must not be underestimated, because they may prove difficult and expensive to resolve. There are four stages required for this analysis:

1. Implementation risk or feasibility
2. Cost to implement and support
3. Effectiveness in increasing control, thereby reducing risk
4. Data classification

Implementation risk is a function of the organization's ability to effectively roll out the technology. In this example, we assume that implementation risk is low and resources are readily available to implement the technology.

Cost to implement and support are paramount to the decision-making process. Both businesses' costs need to be considered. Unfortunately, the cost to support is difficult to quantify and easily overlooked. In this example, resources are available both to implement and to support the technology.

Effectiveness in increasing control is an integral part of the benefit and risk-management analysis. Each security component needs to be considered in order to determine cross-benefits where controls overlap and supplement others. In this example, increased levels of control are understood and supported by management.

Data classification is the foundation for requirements and will help drive the cost-benefit discussions because it captures the value of the information to the underlying business. In this example, the data are considered significant enough to warrant additional security measures to safeguard the data against misuse, theft, and loss.

Each security service can be defined along a continuum, with implementation risk, cost, and data classification all considered. Management can use this chart to plot the security levels required by the system. This example outlines the effectiveness of security services options relative to other protocols or products. Each organization should develop its own continuums and provide guidance to Web developers and application programmers as to the correct uses and standard settings of the security services. For example, for authentication/identification services, in order of minimal service to maximum:

- The minimal level of security is to have no passwords.
- Weak passwords (e.g., easy to guess, shared, poor construction) are better than no passwords but still provide only a minimal level of security.
- Operating system or database level passwords usually allow too much access to the system but can be effectively managed.
- Application passwords are difficult to manage but can be used to restrict data access to a greater degree.
- Role-based access distinguishes users by their need to know to support their job function. Roles are established and users are grouped by their required function.
- Tokens are given to users and provide for two-part authentication. Passwords and tokens are combined for strong authentication.
- Biometrics are means to validate the person claiming to be the user via fingerprints, retina scans, or other unique body function.

For more information on identification and authentication, see Chapters 28 and 29 in this *Handbook*.

30.2.6 Framework Conclusion. Internet e-commerce has changed the way corporations conduct business with their customers, vendors, suppliers, and business units. The B2B and B2C sectors will only continue to grow. Despite security concerns, the acceleration toward increased use of the Internet as a sales, logistics, and marketing channel continues. The challenge for information security professionals' is to keep pace with this change from a security perspective, but not to impede progress. Another equal challenge is that the products that secure the Internet are new and not fully functional or mature. The products will improve, but meanwhile, existing products must be implemented, and later retrofitted, with improved and more secure security services. This changing environment, including the introduction of ever more serious and sophisticated threats, remains difficult to secure.

The processes described in this section will allow the security practitioner to provide business units with a powerful tool to communicate, select, and implement information security services. Three steps were described and demonstrated with two examples. The process supports decision making. Decisions can be made and readily documented to demonstrate cost effectiveness of the security selections. The risk of specific decisions can be discussed and accepted by management. The trade-offs between cost and benefit can be calculated and discussed. Therefore, it becomes critical that alternatives be

reviewed and good decisions made. The processes supporting these decisions need to be efficient and quickly applied. The information security services approach will allow companies to implement security at a practical pace. Services not selected are easily seen. The risk of not selecting specific security services needs to be accepted by management.

30.3 RULES OF ENGAGEMENT. The Web is a rapidly evolving, complex environment. Dealing with customers electronically is a challenge. Web-related security matters raise many sensitive security and privacy issues. Attacks against Websites are not private incidents; they are often public and must always be taken seriously. Correctly differentiating "false alarms" from real attacks is an ongoing challenge with little hope of clean solutions. Distinguishing a popularity spike from a Distributed Denial of Service (DDoS) attack is imprecise and error prone. Indeed, a single traffic surge could be both simultaneously. The Victoria's Secret online lingerie show in February 1998 was popular beyond even the most optimistic expectations of its creators, and the volume of visitors caused severe problems. Obviously, the thousands of people were not attacking the site; they were merely a virtual mob attempting to access the same site at the same time. Similar episodes have occurred when sites were described as interesting on Usenet newsgroups or on social networking sites (e.g., YouTube, Facebook, Twitter, and others). Request tsunamis often occur with little or no warning. Physical mobs are limited by transportation, timing, and costs; virtual mobs are solely limited by how many can attempt to access a resource simultaneously, from locations throughout the world. A recent example was the debut of a limited Missoni product set by Target on September 13, 2011, which generated a demand surge that rendered Target's Website effectively inaccessible (93 percent of requests were refused) for hours.[9,10]

30.3.1 Website-Specific Measures. Protecting a Website means ensuring that the site and its functions are available 24 hours a day, seven days a week, and 365 days a year. It also means ensuring that the information exchanged with the site is accurate and secure.

This section focuses on issues specific to Web interactions with customers as well as supply and distribution chains. Practically speaking, the Website is an important, if not the most important, component of an organization's interface with the outside world.

Website protection lies at the intersection of technology, strategy, operations, customer relations, and business management. Website availability and integrity directly affect the main streams of cash flow and commerce: an organization's customers, production chains, and supply chains.

Availability is the cornerstone of all Web-related strategies. Idle times have become progressively more sparse. Depending on the business and its markets, there may be some periods of lower activity. In the financial trading community, there remain only a few small windows during a 24-hour period when updates and maintenance can be performed. As global business becomes the norm, customers, suppliers, and distributors increasingly expect information, and the ability to effect transactions, at any time of the day or night, even from modest-size enterprises. On the Internet, "nobody knows that you are a dog" also means "nobody knows that you are *not* a large company." The playing field has indeed been leveled, but it was not uniformly raised or lowered. Expectations have increased while capital and operating expenses have dramatically dropped.

Causation is unrelated to impact. The overwhelming majority of Web outages are caused by unglamorous problems. High-profile, deliberate attacks are much less

frequent than equipment and personnel failures. The effect on the business organization is indistinguishable. Having a low profile is no defense against random scanning attack or someone looking for a springboard from which to attack a third party.

Additionally, while an individual firm's assets may not be an attractive target, the firm can still be a useful springboard for an attack against a third party. Such attacks may occur because of a relationship (e.g., vendor or supplier) or may merely be opportunistic attempts at recruitment (e.g., recruiting, or perhaps more appropriately, press-ganging desktop computers for botnets).

External events and their repercussions can also wreak havoc, both directly and indirectly. The September 11, 2001, terrorist attacks that destroyed New York City's World Trade Center complex had worldwide impact, not only on systems and firms located in the destroyed complex. Telecommunications infrastructure was damaged or destroyed, severing Internet links for many organizations. Parts supply and all travel was disrupted when North American airspace was closed for several days. Manhattan was sealed to exits and entries, while within the city itself, and throughout much of the world, normal operations were suspended. The September 11 attacks were extraordinarily disruptive, but security precautions similar to those described throughout this *Handbook* served to ameliorate damage to Web operations and other infrastructure elements of those concerns that had implemented them. Indeed, the existence of the Web and the resulting ability to organize groups without physical presence proved a means to ameliorate the damage from the attacks, even to firms that had a major presence in the World Trade Center. In the period following the attacks on the World Trade Center, Morgan Stanley and other firms with offices in the affected area implemented extensive telecommuting and Web-based interactions first to account for their staffs[11] and then to enable work to continue.[12]

Best practices and scale are important. Some practices, issues, and concerns at first glance appear relevant only to very large organizations, such as Fortune 500 companies. In fact, this is not so. Considering issues in the context of a large organization permits them to appear magnified and in full detail. Smaller organizations are subject to the same issues and concerns but may be able to implement less formal solutions. "Formal" does not necessarily imply written procedures. It may mean that certain computer-related practices, such as modifying production facilities in place, are inherently poor ideas and should be avoided. Very large enterprises might address the problems by having a separate group, with separate equipment, responsible for operating the development environment. Incremental staged deployments are an excellent technique to calibrate risk exposure.

30.3.2 Defining Attacks. Repeated, multiple attempts to connect to a server could be ominous, or they could be nothing more than a customer with a technical problem. Depending on the source, large numbers of failed connects or aborted operations coming from gateway nodes belonging to an organization could represent a problem somewhere in the network, an attack against the server, or a blend of both. It could also represent something no more ominous than a group of users within a locality accessing a Web resource through a firewall.

30.3.3 Defining Protection. There is a difference between protecting Internet-visible assets and protecting Websites. For the most part, Internet-visible assets are not intended for public use. Thus, it is often far easier to anticipate usage volumes and to account for traffic patterns. Websites are subject to the vicissitudes of worldwide public usage. A dramatic surge in traffic could be an attack, or it could be an unexpected

display of the site's URL in a television program or in a relatively unrelated news story. Differentiating between belligerence and popularity is difficult, if not impossible.

Self-protective measures that do not impact customers are always permissible. However, care must be exercised to ensure that the measures are truly impact-free. As an example, some sites, particularly public FTP servers, often require that the Internet protocol (IP) address of the requesting computer have an entry in the inverse domain name system, which maps IP addresses to host names (e.g., node 192.168.0.1 has a PTR [pointer record] 1.0.168.192.in-addr.arpa) (RFC1034, RFC1035; Mockapetris 1987a, 1987b) as opposed to the more widely known domain name system database, which maps host names into IP addresses. It is true that many machines have such entries, but it is also true that many sites, including company networks and many ISPs, do not provide inverse DNS information. Whether this entire population should be excluded from the site is a policy and management decision, not a purely technical decision. Even a minuscule incident rate on a popular Website can be catastrophic, both for the provider and for the naïve end user who has no power to understand or resolve the situation.

30.3.4 Maintaining Privacy. Logging interactions between customers and the Website is also a serious issue. A Website's privacy policy is again a managerial, legal, and customer relations issue with serious overtones. Technical staff needs to be conscious that policies, laws, and other issues may dictate what information may be logged, where it can be stored, and how it may be used. For example, the 1998 Children's Online Privacy Protection Act (COPPA) (15 U.S.C. § 6501 et seq.) makes it illegal to obtain name and address information from children under the age of 13 in the United States. Many firms are party to agreements with third-party organizations such as TRUSTe,[13] governing the use and disclosure of personal information. For more information on legal aspects of protecting privacy, see Chapter 69 in this *Handbook*.

30.3.5 Working with Law Enforcement. Dealing with legal authorities is similarly complicated. Attempts at fraudulent purchases and other similar issues can be addressed using virtually the same procedures that are used with conventional attempts at mail or phone order fraud. Dealing with attacks and similar misuses is more complicated and depends on the organization's policies and procedures, and the legal environment. The status of the Website is also a significant issue. If the server is located at a hosting facility, or is owned and operated by a third party, the situation becomes even more legally complicated. Involving law enforcement in a situation will likely require that investigators have access to the Web servers and supporting network, which may be difficult. Last, there is a question of what information is logged, and under what circumstances. For more information on working with law enforcement, see Chapter 61 in this *Handbook*.

30.3.6 Accepting Losses. No security scheme is foolproof. Incidents will happen. Some reassurance can be taken from the fact that the most common reasons for system compromises in 2001 appear to remain the same as when Clifford Stoll wrote *The Cuckoo's Egg* in 1989. Then and now, poorly secured systems have:

- Management accounts with obvious passwords
- Unprotected system files
- Unpatched known security holes

However, eliminating the simplest and most common ways in which outsiders can compromise Websites does not resolve all problems. The increasing complexity of site content, and of the applications code supporting dynamic sites, means that there is an ongoing design, implementation, testing, and quality assurance challenge. Server-based and server-distributed software (e.g., dynamic www sites) is subject to the same development hazards as other forms of software. Security hazards will slip into a Website despite the best efforts of developers and testers. The acceptance of this reality is an important part of the planning necessary to deal with the inevitable incidents. When it is suspected that a Website, or an individual component, has been compromised, the reaction plans should be activated. The plans required are much the same as those discussed in Chapter 56 in this *Handbook*. The difference is that the reaction plan for a Website has to take into consideration that the group primarily impacted by the plan will be the firm's customers. The primary goal of the reaction plan is to contain the damage. For more information on computer security incident response, see Chapter 56.

30.3.7 Avoiding Overreaction. Severe reactions may create as much, if not more, damage than the actual attack. The reaction plan must identify the decision-making authority and the guidelines to allow effective decisions to be made. This is particularly true of major sites, where attacks are likely to occur on a regular basis. Methods to determine the point at which the Website must be taken offline to prevent further damage need to be determined in advance.

Offline clones should be prepared to enable fast responses. When under time pressure during an emergency, the last thing one wants to do is building production system images.

In summary, when protecting Websites and customers, defensive actions are almost always permissible and offensive actions of any kind are almost always impermissible. Defensive actions that are transparent to the customer are best of all.

30.3.8 Appropriate Responses to Attacks. Long before the advent of the computer, before the development of instant communications, international law recognized that firing upon a naval vessel was an act of war. Captains of naval vessels were given standing orders summarized as *fire if fired upon*. In areas without readily accessible police protection, the right of citizens to defend themselves is generally recognized by most legal authorities. Within the body of international law, such formal standards of conduct for military forces are known as *rules of engagement*, a concept with global utility.

In cyberspace, it is tempting to jettison the standards of the real world. It is easy to imagine oneself master of one's own piece of cyberspace, without connection to real-world laws and limitations on behavior. However, information technology (IT) personnel do not have the legal standing of ships' captains with no communications to the outside world. Some argue that "fire if fired upon" is an acceptable standard for online behavior. Such an approach does not take into account the legal and ethical issues surrounding response strategies and tactics. It may be unsatisfying, but the only response relatively free of legal consequences (criminal, civil, and regulatory) is "Defensive Action Only." Blocking address ranges is defensive; firing back crosses the line into offense.

Any particular security incident has a range of potential responses. Which response is appropriate depends on the enterprise and its political, legal, and business environment. Acceptability of response is also a management issue as well as potentially a political

issue. Determining what responses are acceptable in different situations requires input from management on policy, from legal counsel on legality, from public relations on public perceptions, and from technical staff on technical feasibility. Depending on the organization, it also may be necessary to involve unions and other parties in the negotiation of what constitutes appropriate responses.

What is acceptable or appropriate in one area is not necessarily acceptable or appropriate in another. Often the national security arena has lower standards of proof than would be acceptable in normal business litigation. In U.S. civil courts, cases are decided upon a preponderance of evidence. Standards of proof acceptable in civil litigation are not conclusive when a criminal case is being tried, where guilt must be established beyond a reasonable doubt.

Gambits or responses that are perfectly legal in a national security environment may be completely illegal and recklessly irresponsible in the private sector, exposing the organization to significant legal liability.

Rules of etiquette and behavior are similarly complex. The rights of prison inmates in the United States remain significant, even though they are subject to rules and regulations substantially more restrictive than for the general population. Security measures, as well, must be appropriate for the persons and situations to which they are applied.

30.3.9 Counter-Battery. Some suggest that the correct response to a perceived attack is to implement the cyberspace equivalent of *counter-battery,* that is, targeting the artillery that has just fired upon you. However, counter-battery tactics, when used as a defensive measure against Internet attacks, will be perceived, technically and legally, as an attack like any other.

Counter-battery tactics may be emotionally satisfying but are prone to both error and collateral damage. Counter-battery can be effective only when the malefactor is correctly identified and the effects of the reciprocal attack are limited to the malefactor. If third parties are harmed in any way, then the retaliatory action becomes an attack in and of itself. One of the more celebrated counter-battery attacks gone awry was the 1994 case when two lawyers from Phoenix, Arizona, spammed over 5,000 Usenet newsgroups to give unsolicited information on a U.S. Immigration and Naturalization Service lottery for 55,000 green cards (immigration permits). The resulting retaliation—waves of email protests—against the malefactors flooded their Internet service provider (ISP) and caused at least one server to crash, resulting in a denial of service to all the other, innocent, customers of the ISP.[14]

30.3.10 Hold Harmless. Operational policies for Internet operations must adopt a hold harmless position vis-à-vis good faith actions within policy. Dealing with an Internet crisis often requires fast reactions. If employees act in good faith, in accordance with their responsibilities, and within documented procedures, those making the decisions should be immune from reprisal. Punishment for actions undertaken in good faith is counterproductive. If the procedures are wrong, improve the rules and procedures; do not blame the employees for following established policies. Disciplinary actions are manifestly inappropriate in such circumstances.

30.4 RISK ANALYSIS. As noted earlier in this chapter, protecting an organization's Websites depends on an accurate, rational assessment of the risks. Developing effective strategies and tactics to ensure site availability and integrity requires that all potential risks be examined in turn.

Unrestricted commercial activity has been permitted on the Internet since 1991. Since then, enterprises large and small have increasingly integrated Web access into their second-to-second operations. The risks inherent in a particular configuration and strategy are dependent on many factors, including the scale of the enterprise and the relative importance of the Web-based entity within the enterprise. Virtually all high-visibility Websites (e.g., Yahoo!, America Online, cnn.com, Amazon.com, and eBay) have experienced significant outages at various times.

The more significant the organization's Web component, the more critical is availability and integrity. Large, traditional firms with relatively small Web components can tolerate major interruptions with little damage. Firms large or small that rely on the Web for much of their business must pay greater attention to their Web presences, because a serious outage can quickly escalate into financial or public relations catastrophe.

For more details of risk analysis and management, see Chapters 62 and 67 in this *Handbook*.

30.4.1 Business Loss. Business losses fall into several categories, any of which can occur in conjunction with an organization's Web presence. In the context of this chapter, customers are both outsiders accessing the Internet presence and insiders accessing intranet applications. In practice, insiders using intranet-hosted applications pose the same challenges as the outside users.

30.4.2 PR Image. The Website is the organization's public face 24/7/365. This ongoing presence is a benefit, making the firm visible at all times, but the site's high public profile also makes it a prime target.

Government sites in the United States and abroad have often been the targets of attacks. In January 2000, "Thomas," the Website of the U.S. Congress, was defaced. Earlier, in 1996, the Website of the U.S. Department of Justice was vandalized. Sites belonging to the Japanese, U.K., and Mexican governments also have been vandalized.

Attacks against Websites continue to be an ever-increasing problem. On March 21, 2013, a series of malware infections crippled broadcasters and financial institutions in South Korea during a time of increasing tensions over North Korean nuclear and missile testing, including a renunciation by North Korea of the 1953 armistice suspending the Korean War.[15]

These incidents continue, with hackers on both sides of various issues attacking the other side's Web presence. While no conclusive evidence of the involvement of nation states in such activities has become generally known, it is inevitable. Surges of such activity have coincided with major public events, including the 2001 U.S.–China incident involving an aerial collision between military aircraft, the Afghan and Iraqi operations, and the Israeli–Lebanese war of 2006, the Russian–Georgian conflict,[16] and Estonia.[17] Such activity has not been limited to the national security arena, however. Many sites were defaced during the incidents following publication of cartoons in a Danish newspaper that some viewed as defaming the prophet Muhammad.

In some cases, companies have suffered collateral damage from hacking contests, where hackers prove their mettle by defacing as many sites as they can.[18] The fact that there is no direct motive or animus toward the company is not relevant; the damage has still been done.

In the corporate world, company Websites have been the target of attacks intended to defame the corporation for real or imagined slights. Some such episodes have been reported in the news media, whereas others have not been the subjects of extensive reporting. The scale and newsworthiness of the episode is unimportant; the damage

done to the targeted organization is the true measure. An unreported incident that is the initiating event in a business failure is more damaging to the affected parties than a seemingly more significant outage with less severe consequences.

Other cybervandals (e.g., sadmind/IIS) have used address scanners to target randomly selected machines. Obscurity is not a defense against address-scanner attacks.

30.4.3 Loss of Customers/Business. Internet customers are highly mobile. Website problems quickly translate into permanently lost customers. The reason for the outage is immaterial; the fact that there was a problem is often sufficient to provoke customer flight.

In most areas, there is competitive overlap. Using the overnight shipping business as an example, in most U.S. metropolitan areas there are (in alphabetical order) Federal Express, United Parcel Service, and the United States Postal Service. All of the firms offer Web-based shipment tracking, a highly popular service. Problems or difficulties with shipment tracking will quickly lead to a loss of business in favor of a different company with easier tracking.

30.4.4 Interruptions. Increasingly, modern enterprises are being constructed around ubiquitous 24/7/365 information systems, most often with Websites playing a major role. In this environment, interruptions of any kind are catastrophic.

Production. The past 20 years have seen a streamlining of production processes in all areas of endeavor. Twenty years ago, it was common for facilities to have multiday supplies of components on-hand in inventory. Today, *zero latency* or *just-in-time* (JIT) environments are common, permitting large facilities to operate with minimal inventory. Fiscally, zero latency environments may be optimally cost efficient, yet the paradigm leaves little margin for disruptions of the supporting logistical chain. This chain is sometimes fragile, and subject to disruption by any number of hazards.

Supply Chain. Increasingly, it is common for Web-based sites to be an integral part of the supply chain. Firms may encourage their vendors to use a Web-based portal to gain access to the vendor side of the purchasing system. XML[19]-based gateways and *service-oriented architecture* (SOA) approaches, together with other Web technologies, are used to arrange for and manage the flow of raw materials and components required to support production processes. The same streamlining that speeds information between supplier and manufacturer also provides a potential for serious mischief and liability.

Delivery Chain. Web-based sites, both internal and external, human readable and intersystem, have become the vehicle of choice for tracking the status of orders and shipments. Increasingly system-system links, implemented as electronic-documents (e.g., XML or JSON) have become the backbone of many enterprises' delivery chain management and inquiry systems.

Information Delivery. Banks, brokerages, utilities, and municipalities are increasingly turning to the Web as a convenient, low-cost method for managing their relationships with consumers. Firms are also supporting downloading records of transactions and other relationship information in formats required by personal database programs and organizers. These outputs, in turn, are often used as inputs to other processes, which then generate other transactions. Not surprisingly, as time passes, more and more people and businesses depend on the availability of information on demand. Today's Web-based customers presume that information is accessible wherever they can use a smartphone or tablet. This is reminiscent of usage patterns of automatic teller machines in an earlier decade, which allowed people to base their plans on access

to teller machines, often making multiple $20 withdrawals instead of cashing a $200 check weekly.

30.4.5 Proactive versus Reactive Threats. Some threats and hazards can be addressed proactively, whereas others are inherently reactive. When strategies and tactics are developed to protect a Web presence, the strategies and tactics themselves can induce availability problems.

As an example, consider the common strategy of having multiple name servers responsible for providing the translation of domain names to IP addresses. It is required before a domain name (properly referred to as a Domain Name System [DNS] zone) can be entered into the root-level name servers, that at least two name servers be identified to process name resolution requests. Name servers are a prime example of resources that should be geographically diverse.

Updating DNS zones requires care. If an update is performed improperly, then the resources referenced via the symbolic DNS names will become unresolvable, regardless of the actual state of the Web server(s) and related infrastructure. The risk calculus involving DNS names is further complicated by the common, efficient, and appropriate practice of designating ISP name servers as the primary mechanism for the resolution of domain names. In short, name translation provides a good example of the possible risks that can affect a Web presence.

30.4.6 Threat and Hazard Assessment. Some threats are universal. Others are specific to an individual environment. The most devastating and severe threats are those that simultaneously affect large areas or populations, where efforts to repair damage and correct the problem are hampered by the scale of the problem. Differences in perspective can impede effective response. What seems a minor outage to a carrier may be a business-threatening outage to a customer.

On a basic level, threats can be divided into several categories. The first is between deliberate acts and accidents. Deliberate acts comprise actions done with the intent to damage the Website or its infrastructure. Accidents include natural phenomena (acts of God) and clumsiness, carelessness, and unconsidered consequences (acts of clod).

Broadly put, a deliberate act is one whose goal is to impair the system. Deliberate acts come in a broad spectrum of skill and intent. For the purpose of risk analysis and planning, deliberate acts against infrastructure providers can often appear to be acts of God. To an organization running a Website, an employee attack against a telephone carrier appears simply as one more service interruption of unknown origin.

No enterprise or agency should consider itself an unlikely target. Past high-profile incidents have targeted the FBI (May 26 and 27, 1999), major political parties, and interest groups in the United States. On the consumer level, numerous digital subscriber line (DSL)-connected home systems have been targeted for subversion as preparation for the launching of DDoS attacks. In 2007, several investigations resulted in the arrest of several individuals for running so-called botnets (ensembles of compromised computers). These networks numbered hundreds of thousands of machines.[20] Recently, Aramco was victimized by a major attack[21]; as were major Korean financial institutions.[22] The potential for such networks to be used for mischief cannot be underestimated.

30.5 OPERATIONAL REQUIREMENTS. Internet-visible systems are those with any connection to the worldwide Internet. It is tempting to consider protecting

Internet-visible systems as a purely technical issue. However, technical and business issues are inseparable in today's risk management. For example, as noted earlier in this chapter, the degree to which systems should be exposed to the Internet is fundamentally a business risk-management issue. Protection technologies and the policies behind the protection can be discussed only after the business risk questions have been considered and decided, setting the context for the technical discussions. In turn, business risk-management evaluation (see Chapter 62) must include a full awareness of all of the technical risks. Ironically, nontechnical business managers can accurately assess the degree of business risk only after the technical risks have been fully exposed.

Additional business and technical risks result from outsourcing. Today, many enterprises include equipment owned, maintained, and managed by third parties. Some of this equipment resides on the organization's own premises and other equipment resides offsite: for example, at application service provider facilities.

Protecting a Website begins with the initial selection and configuration of the equipment and its supporting elements, and continues throughout its life. In general, care and proactive consideration of the availability and security aspects of the site from the beginning will reduce costs and operational problems. Although virtually impossible to achieve, the goal is to design and implement an automatic system, with a configuration whose architecture and implementation operates even in the face of problems, with minimal customer impact.

That is not to say that a Website can operate without supervision. Ongoing, proactive monitoring is critical to ensuring the secure operation of the site. Redundancy only reduces the need for real-time response by bypassing a problem temporarily; it does not eliminate the underlying cause. The initial failure must be detected, isolated, and corrected as soon as possible, albeit on a more schedulable basis. Otherwise, the system will operate in its successively degraded redundancy modes until the last redundant component fails, at which time the system will fail completely.

30.5.1 Ubiquitous Internet Protocol Networking. Business has been dealing with the security of internets (i.e., interconnected networks) since the advent of internetworking in the late 1960s. However, the ubiquity of Transmission Control Protocol/Internet Protocol (TCP/IP) networks and of the public Internet has exposed much more equipment to attack than in the days of closed corporate networks. In addition, a much wider range of equipment, such as voice telephones based on voice-over IP (VoIP), fax machines, copiers, and even soft drink dispensers, are now network accessible.

IP connectivity has been a great boon to productivity and ease of use, but it has not been without a darker side. Network accessibility also has created unprecedented opportunities for improper, unauthorized access to networked resources and other mischief. It is not uncommon to experience probes and break-in attempts within hours or even minutes of unannounced connection to the global Internet.

Protecting Internet-visible assets is inherently a conflict between ease of access and security. The safest systems are those unconnected to the outside world. Similarly, the easiest to use systems are those that have no perceivable restrictions on use. Adjusting the limits on user activities and access must balance conflicting requirements. As an example, many networks are managed *in-band*, meaning that the switches, routers, firewalls, and other elements of the network infrastructure are managed using the actual network as the connection medium. If the management interfaces were not managed over properly encrypted connections, management passwords would be visible on the

network. An outsider or unauthorized insider may thus monitor the network, gaining sufficient information to paralyze the network. This can result in severe damage and liability to the organization.

30.5.2 Internal Partitions. Complex corporate environments can often be secured effectively by dividing the organization into a variety of interrelated and nested security domains, each with its own legal, technical, and cultural requirements. For example, there are specific legal requirements for medical records (see Chapter 71 in this *Handbook*) and for privacy protection (see Chapter 69). Partners and suppliers, as well as consultants, contractors, and customers, often need two-way access to their corporate data and facilities. These diverse requirements mean that a single corporate firewall is often insufficient. Different domains within the organization will often require their own firewalls and security policies. Keeping track of the multitude of data types, protection and access requirements, and different legal jurisdictions and regulations makes for previously unheard-of degrees of complexity. It is simply not possible to configure a complete set of rules to implement the multiple, different security requirements without creating security gaps or contradictions.

Much has been written about so-called *Advanced Persistent Threats*.[23] While the terminology is relatively new, the concept is not, having been mentioned in earlier editions of this handbook.[24] The fundamental underlying change with advanced threats is the degree of targeting and the attacker's intention to remain covert. In these respects a useful historic precedent is involvement in intelligence gathering by submarine forces. During World War II, submarines were first explicitly tasked with intelligence gathering. The orders given can be summarized as "get the information and/or photographs; avoid detection at all costs." The translation is straightforward: No matter how attractive, do not announce your presence by making an attack; your purpose is to stay covert and extract the desired information. This is precisely the model of today's targeted, information extraction attacks.

Damage control is another property of a network with internal partitions. A system compromised by an undetected malware component will be limited in its ability to spread the contagion beyond its own compartment.

30.5.3 Critical Availability. Networks are often critical for second-to-second operations. The side effects of ill-considered countermeasures may be worse than the damage from an actual attack. For example, shutting down the network or even part of it, for maintenance or repair can wreak more havoc than penetration by a malicious hacker. Damage control, often starting with compartmentalization, is a critical part of limiting the impact of an attack.

30.5.4 Accessibility. Users must be involved in the evolution of rules and procedures. Today, it is still not unheard of for a university faculty to take the position that any degree of security will undermine the very nature of their community, compromising their ability to perform research and inquiries. This extreme position persists despite the attention of the mass media, the justified beliefs of the technical community, and documented evidence that lack of protection of any Internet-connected system undermines the safety of the entire connected community.

Connecting a previously isolated computer system or network to the global Internet creates a communications pathway to every corner of the world. Customers, partners, and employees can obtain information, send messages, place orders, and otherwise interact 24 hours a day, seven days a week, 365 days a year, from literally anywhere on

or near Earth. The Space Shuttle had Internet access; as does the International Space Station. Under these circumstances, the possibilities for attack or inadvertent misuse are limitless.

Despite the vast increase in connectivity, some businesses and individuals do not need extensive access to the global Internet for their day-to-day activities, although they may resent being excluded. The case for universal access is therefore a question of business policy and political considerations.

30.5.5 Applications Design. Protecting a Website begins with the most basic steps. First, a site processing confidential information should always support the Secure Hypertext Transfer Protocol (HTTPS), typically using Transmission Control Protocol (TCP) port 443. Properly supporting HTTPS requires the presence of an appropriate digital certificate (see Chapter 37).

When the security requirements are uncertain, the site design should err on the side of using HTTPS for communications. Eavesdropping on the Internet is an ever-present hazard. There have been cases of routing database manipulation and state-sponsored surveillance. The Hannaford and TJ Maxx episodes described earlier in this chapter both involved the application of what can only be described as classic signals intelligence (SIGINT) techniques. Even though details on such attacks are often not made public, the voluminous historical literature on communications intelligence (COMINT) and SIGINT from the World War II period makes one point abundantly clear: Encryption of all potentially sensitive traffic is the only way to protect information.

Encryption also should be used within an organization, possibly with a different digital certificate, for sensitive internal communications and transactions. Earlier, it was noted that organizations are not monolithic security domains. This is nowhere more accurate than when dealing with human resources, employee evaluations, compensation, benefits, and other sensitive employee information. There are positive requirements that this information be safeguarded, but few organizations have truly secured their internal networks against internal monitoring. It is far safer to route all such communications through securely encrypted channels. Such measures also demonstrate a good faith effort to ensure the privacy and confidentiality of sensitive information. Such efforts are an important element of an organization's defense when an incident occurs. Upfront care can significantly reduce legal liability.

Failure to heed these hazards has been at the root of several incidents, including the TJ Maxx and Hannaford affairs.

It is also important to avoid providing all of the authentication information on a single page or, for that matter, in a sequence of pages. When parts of information are suppressed, as, for example, portions of a credit card or account number, the division between suppressed and displayed portions should be maintained. Displaying all of the information, even if it is on different screens, is an invitation to a security breach.

30.5.6 Provisioning. Although today's hardware has unprecedented reliability, any failure of hardware between the customer and the data center will impair an enterprise's Web presence. For a highly available, front-line Website, the effective requirement is a minimum of two diversely located facilities, each with a minimum of two servers. This is not necessarily an expensive proposition. Fairly powerful Web servers can be purchased for less than $5,000, so the total hardware expenditure for four servers is reasonable, substantially less than the annual cost of a single technician. Cloud-computing, the availability of data center-housed virtual servers as a service (e.g.,

Amazon, Verizon, RackSpace, Google) presents another option with low entry costs. In most cases, the cost of the extra hardware is more than offset by the business cost of downtime, which can sometimes exceed the total cost of the duplicative hardware by a factor as much as 100, in a single episode.[25]

Duplicate hardware and geographic diversity ensure constant customer access to some degree of functionality. The degree of functionality that must be maintained depends on the market and the customers. Financial firms supporting online stock trading have different operational, regulatory, and legal requirements than supermarkets. The key is matching the support level to the activities. Some degree of planned degradation is generally acceptable. Total unavailability is not an option.

Remote hosting services generally referred to as "cloud computing" do not eliminate risks. While renting time and other resources from a provider does reduce capital and operating costs, there is a corresponding loss of visibility and control. A failure at a remote vendor operated facility can take down a critical service. Depending on your applications, industry, regulatory, and legal environments, the risks of cloud computing may be quantifiable, or they may be immeasurable. Difficulties in vendor relationships can give rise to "hostage data" situations. From a security and information assurance perspective, planning in advance to deal with changes in cloud vendors should be part of planning.

There are also unresolved legal issues surrounding the use of shared computing facilities operated by third parties. If a search warrant is served on an internal facility, the organization will clearly be aware of what has happened. If a shared facilities vendor is served with a warrant, it is an open question whether the vendor has the desire (or capability) to present the defenses clearly available with data stored within the company.

30.5.7 Restrictions. All Web servers should be located behind a firewall in a DMZ, as discussed in Section 30.6.4. Incoming and outgoing services should be restricted using protocols such as HTTP, HTTPS, and Internet Control Message Protocol (ICMP). For troubleshooting purposes, it is desirable to implement ICMP, which is used by PING, an echo requester, as a way to check connectivity. All unused ports should be disabled. Furthermore, the disabled ports should be blocked by the firewalls separating the DMZ from the outside world.

Customer information should, to the extent possible, be stored on systems separate from the systems actually providing Web serving. Many security episodes have exploited file protection errors on Web servers, in order to access databases directly. Segregating customer data on separate machines, and ensuring that the only way to access customer data is through the documented pathways, is likely severely to impede improper attempts to access and modify information.

These safeguards are especially important for high-security information such as credit card numbers. The number of incidents in which malefactors have downloaded credit card numbers directly from Website is an indication of the importance of such precautions. The systems actually storing the sensitive information should never be accessible from the public Internet. Storage and processing of payment card information is also subject to industry guidelines, particularly the PCI DSS.

The TJX case, which came to public attention in the beginning of 2007, was one of a series of large-scale compromises of electronically stored information on back-office and e-commerce systems. Most notably, the TJX case appears to have started with an insufficiently secured corporate network and the associated back-office systems, not

a Website penetration. This breach escalated into a security breach of corporate data systems. It is reported on the TJX Website that at least 45.7 million credit cards were compromised (original reports in *USA Today* and other publications cite 94 million credit card numbers as being compromised). On November 30, 2007, it was reported that TJX, the parent organization of stores including TJ Maxx and Marshall's, had agreed to settle bank claims related to VISA cards for US$ 40.9M.[26] Also, as of March 2008, a class action suit on behalf of customers was in process of being settled. It called for fees to be paid to each of the plaintiffs for credit monitoring and identity theft, and for new driver's licenses, as well as $6.5 million to plaintiff's counsel. In spite of several such high-profile attacks, the lessons appear not to have been learned. As recently as March 2008, Hannaford Bros. Co. reported that 4.2 million credit card numbers had been exposed, with about 1,800 cases of fraudulent usage reported as of that date. An incident involving the network supporting Sony's PlayStation 3[27] disclosed personal information relating to an estimated 70 million subscribers. Clearly, more stringent controls, care, and safeguards are called for.

30.5.8 Multiple Security Domains. The front-line Web servers, and the database servers supporting their activities, comprise two different security domains.

The Web servers, as noted previously, need to be globally accessible via HTTP, HTTPS, and ICMP (while not needed for actual operations, ICMP is often needed for troubleshooting connectivity problems by both customers and providers). In turn, they need to access the application or database servers, and *only* those servers. In a production environment, it is preferable that application or database servers interact with the Web servers using a dedicated, restricted-use protocol. Properly implemented, such a restriction prevents a hijacked Web server from exploiting its access to the application or database server.

These second-tier servers should be in a security domain separated by restrictive firewalls from the externally accessible front-line Web servers. This seems like a significant expenditure, but it is often less expensive and lower risk than the costs associated with correcting a single significant incident.

30.5.9 What Needs to Be Exposed? Publicly accessible Websites need publicly accessible Web servers to perform their functions. The challenge is to provide the desired services to the public without simultaneously providing access that can be leveraged in unauthorized ways to subvert the site. A penetration incident may lead to significant financial losses, embarrassment, and financial and (in some cases) criminal liability.

Generally speaking, Web servers should not be connected to the public Internet (unless it is possible to verify that all other ports are closed to traffic). All connections to the public network should be made through a firewall system, with the firewall configured to pass only Web-related traffic to those hosts. A DMZ guarded by a firewall often creates a more easily controlled security gateway.

Many sites will opt to place externally visible Web servers in a security compartment of their own, on a separate port of the firewall (if not a totally separate firewall), using a separate DMZ from other publicly accessible resources. These precautions may seem excessive, but having improperly secured systems can lead to security breaches that are extremely difficult to correct and can lead to extended downtime while the problems are analyzed and remedied. In this case, an ounce of prevention is worth substantially more than a pound of cure.

30.5.9.1 Exposed Systems. Exposed systems are inherently a security hazard. Systems without access from or to the public network cannot be compromised from the public network. Only systems that absolutely need to be publicly accessible should be so configured. Minimizing the number of exposed systems is generally desirable, but this is best considered in terms of machine roles rather than actual counts of systems. Increasing the load on each publicly accessible server by increasing the size of the server, thus increasing the amount of capacity impacted by a single system outage, is not a benefit. However, this must be balanced against the new trend toward server virtualization, which does not increase the size of a server, but which increases its utilization, in order to lower costs and improve reliability.

The introduction of SSL-based VPN implementations and services supporting remote access via secure HTTPS connections (e.g., gotomypc.com) create an entirely new class of security hazard. SSL and other encrypted techniques have been used to exfiltrate data from security penetrations. While encrypted connections may be difficult or impossible to observe, connections to unexpected locations are a signal of possible security compromise.

30.5.9.2 Hidden Subnets. The servers directly supporting the Website need to be accessed by the outside world, and thus must generally have normal, externally visible Internet addresses. Alternatively, traffic directors, firewalls, or similar devices may map an external port number to an internal address and port number. However, in most cases the other systems supporting the Web servers generally have no legitimate need for unrestricted access from or to the public network.

The safest address assignments for supporting, non–outside-visible systems are the IPv4 addresses allocated for use in private internets[28] and their equivalent IPv6 addresses.[29] Needless to say, these systems should be in a security compartment separated from the publicly accessible Web servers, and that compartment should be isolated from the publicly accessible compartment with a very restrictive firewall.

30.5.10 Access Controls. Publicly accessible systems are both the focus of an organization's security efforts and the primary target of attempts to compromise that security. The number of individuals authorized to make changes to the systems and the ways in which changes may be made need to be carefully controlled, reported, and monitored. The cleared individuals should use individual accounts, and the access to sensitive functions through such accounts should be immediately invalidated if the individual's access is no longer authorized or if the employee ceases employment, is under investigation for impropriety, or other reason. The most efficient and effective way to implement these requirements is role-based access control, which leverages the underlying operating system protection model to enforce file security.[30] The key aspect of role-based protection is that file access is protected by role, not identifier. It is thus possible to remove a user's access to production files with a single command.

Access controls are also not merely a question of traditional role-based control. Increasingly, applications are directly accessed by customers. Such direct access requires more precautions and care than that provided by customer service personnel. It is often appropriate for customer service personnel to have access to most accounts; it is just as inappropriate to give corresponding unfettered access to end customers, who should only have access to their specific account. Applications-level attacks (e.g., SQL injection) become increasingly hazardous.

The transition to Web browser and mobile device–based interfaces needs additional security precautions. Remote clients are subject to compromise and reverse-engineering. As an example, using plaintext, client-supplied information to control access (e.g., cookies or URL-encoded) can easily give rise to an SQL injection (or equivalent) attack, merely by editing the unencoded information. Switching to proper variable, session-based numbering scheme would make tampering far less successful (indeed, the greater number of failures resulting from using an such an opaque number would be easily detectable by applications and security logging systems).

The challenge to such systems is that the safeguards must be implemented within the application. Such safeguards are needed to ensure applications security.

SQL injection attacks of this type have occurred to major organizations. In 2011, details of 200,000 credit card accounts with Citibank were compromised in this way.[31]

For more information on access controls, see Chapters 28 and 29 in this *Handbook*.

30.5.11 Site Maintenance. Maintaining and updating the Website requires great care. The immediate nature of the Web makes it possible for a single-character error in a major enterprise-level application to cause hundreds of thousands of dollars of damage within moments. Web servers need to be treated with respect; the entire enterprise is riding on the electronic image projected by the server. Cybervandalism, which most commonly consists of defacing the home page of a well-known site, requires unauthorized updating of the files comprising the site. In 2000 alone, well-known public and private entities including the FBI, OPEC, World Trade Organization, and NASA, as well as educational institutions including the University of Limerick, have been harmed in this way. These attacks continue to be a danger, both in terms of damage to the organization's image and covertly, as using the WWW site as a launching pad for Web-based exploits.

The Web is inherently a highly leveraged environment. Small changes in the content of a single page percolate throughout the Web in a matter of minutes. Information disseminates easily and quickly at low cost. Leverage helps tremendously when things go right; when things go badly, leverage dramatically compounds the damage. For more information on change control, see Chapters 38, 39, and 40 in this *Handbook*.

30.5.12 Maintaining Site Integrity. Every Website and its servers is a target. Antagonists can be students, activists, terrorists, disgruntled former employees, or unhappy customers as well as state actors. Because Websites are an enterprise's most public face, they represent extremely desirable targets.

Maintaining integrity requires that updates and changes to the site be done in a disciplined manner. Write access to the site must be restricted, and those authorized must use secure methods to access the Web servers. The majority of reported incidents appear to be the result of weak security in the update process. For example, unsecured FTP access from the general Internet is a poor practice. Safer mechanisms include:

- FTP from a specific node within the inner firewall
- KERMIT on a directly wired port or over a secured connection
- Logins and file transfers (e.g., SFTP) over secure authenticated connections via an underlying SSH transport
- Physical media transfers

Most of the technologies do not inherently require that an onsite individual perform server updates, which would preclude remote maintenance. It does mean that in order to get to a machine from which an update can be performed, it is necessary to come through a virtual private network with point-to-point tunneling protocol or Layer2 tunneling protocol (VPN PPTP/L2TP) authenticated by at least one of the secure gateways. For more information on VPNs, see Chapter 32 in this *Handbook*.

30.6 TECHNICAL ISSUES. There are many technical issues involved in protecting Internet-accessible resources. The technologies used to protect network assets include routers, firewalls, proxy servers, redundancy, and dispersion. When properly designed and implemented, security measures produce a positive feedback loop, where improvements in network security and robustness are self-reinforcing. Each improvement makes other improvements possible and more effective.

30.6.1 Inside/Outside. Some visions of the future include a utopian world where everything is directly accessible from anywhere without effort, and with only beneficial results. The original Internet operated on this basis, until a number of incidents (including the 1988 Morris worm) caused people to rethink the perceived inherent peacefulness of the networked world. It is a trend that has only accelerated, as the total number of systems grows ever larger.

The architecture and design of protective measures for a network depend on differentiating inside trustable systems from outside untrustworthy systems. This is equally true for intranets, the Internet, or an *extranet,* so called to distinguish private interconnections of networks from the public Internet. Unfortunately, trust is not a black-and-white issue. A system may be trustworthy from one perspective and untrustworthy from another, thus complicating the security design. In addition, the vast majority of inappropriate computer use is thought to be done by those with legitimate access to some aspect of the system and its data. Alternatively, it is not possible to be strong everywhere. Trusted but compromised systems are a significant danger.

Basic connectivity configuration is one of those few areas that are purely technical, without a business risk element. One of the most obvious elements involves the tables implemented in routers connecting the enterprise to the public carrier–supplied IP connection. The table rules must prevent *IP spoofing*, which is the misrepresentation of IP packet origins. This is also true when originators are within the organization.

There are three basic rules for preventing IP spoofing applicable to all properly configured networks:

1. Packets entering the network from the outside should never have *originator* addresses within the target network.
2. Packets leaving a network and going to the public network must have *originator* addresses within the originating network.
3. Packets leaving a network and going to the public network must not have *destination* addresses within the originating network.

An exception to these rules is in the use of stealth internal networks, those whose internal addresses correspond to legal external addresses.[32]

A corollary to these rules is that packets with originator or destination addresses in the most local intranet addresses range[33] or dynamic IP addresses,[34] should never be

permitted to enter or leave an internal network. Nested address spaces may deliberately create aliased address spaces.[35,36]

30.6.2 Hidden Subnets. Firewalls funnel network traffic through one or more choke points, concentrating the security task in a small number of systems. The reasoning behind such concentration is that the likelihood of security breaches of the entire network rises rapidly with the number of independent access points.

Firewalls and proxy servers (see Chapter 26 in this *Handbook*) are effective only in topologies where the firewall filters all traffic between the protected systems and the less-trusted world outside its perimeter. If the protected systems can in any way be accessed without going through the firewall, then the firewall itself has been rendered irrelevant. Security audits often uncover systems that violate this policy; relying on administrative sanctions to preclude such holes in the security perimeter generally does not work. The rule of international diplomacy, "Trust but verify," applies.

The simplest solution to this problem is the use of RFC 1918[37] addresses within protected networks.[38] RFC1918 provides a range of IPv4 addresses guaranteed by the Internet Assigned Numbers Authority (IANA) never to occur in the normal, public Internet. The address ranges used for dynamic IP address assignment have similar properties.[39] IPv6 networks are similarly governed by RFC 4193.

Filtering these addresses on both inbound and outbound data streams is straightforward and highly effective at stopping a wide range of attacks.[40] Requiring the use of such addresses, and prohibiting the internal use of externally valid addresses, goes a long way toward preventing the use of unauthorized protocols and connections.

30.6.3 What Need Be Exposed? A security implementation starts with an analysis of the enterprise's mission, needs, and requirements. All designs and implementations are a compromise between absolute security, achieved only in a powered-down or disconnected system, and total openness, in which a system is completely open to any and all access from the outside.

Although in most cases communications must be enabled between the enterprise and the outside world for normal business functions, total disconnection, known as an *air gap*, is sometimes both needed and appropriate between specific components. Industrial real-time control systems, life-critical systems, and systems with high requirements for confidentiality remain appropriate candidates for total air gaps.

Often, systems that do not need to receive information from outside sources must publish statistical or other information to less secure systems. This requirement can often be satisfied through the use of limited functionality links such as media exchange, or through tightly controlled one-way transfers. These mechanisms can be implemented with IP-related technologies or with more limited technologies, including KERMIT,[41] UUCP, or vendor-specific solutions.

In other cases, restrictions reflect policies for permitted use and access rather than for protection against outside attack. For example, it is reasonable and appropriate for a public library to limit access to HTTP and related protocols while prohibiting access to such facilities as FTP and TELNET. Remote access using the Secured Shell Protocol (SSH) is extremely useful, as it allows for two-way authentication (both client and server can be authenticated against public/private keys). There is a collateral issue of what provisions should be made for systems within the perimeter to connect to outside networks using various protocols. PPTP and L2TP are the most well-known. Whether these methods of accessing outside networks should be permitted or not is

a management question of no small import, equivalent to totally suppressing outside communications.

From a security standpoint, tunnels to the outside world are a perfect storm, potentially permitting unlimited, nonmonitored communications to the outside. From the standpoint of enabling business operations, such access may be a practical necessity. Vendors, contractors, suppliers, and customers all need access to internal systems at their respective organizations, and such access will invariably require the use of a VPN tunnel. Perhaps the best solution is to reverse the traditional inside/outside dichotomy, making the LAN an untrusted network. The LAN then becomes a universal dial tone. Access to internal corporate systems can then be secured separately using VPN technology.

The advent of SSL/HTTP-based tunnels presents another challenge. gotomypc.com offers such a service, and individuals have used it to circumvent firewalls and implemented ad hoc remote access. The cost is nominal, often a small fraction of the monthly out-of-pocket outlay for an individual's daily commute to the office. The solution of blocking connections to the IP addresses assigned to such a provider is an often recommended, yet inherently flawed, prophylactic. Blocking one or more such services does nothing to secure the network against an as-yet-unidentified service of this type, such as that hosted on a home server. Limiting the duration of SSL connections is also merely a speed bump to such schemes. These remain a challenge, as does the advent of directly usable SSL/HTTP tunnels communicating using the standard HTTP (TCP Port 80) and HTTPS (TCP Port 443) ports.

In these cases, firewalls are a reasonable solution, so long as their limitations are recognized. For example, firewalls do not prevent mobile code such as Active-X and JAVA from working around prohibited functions. Network attacks using code supplied to browsers for local execution within the local network context have been published. These attacks can persist across time, and are one of the attack classes within the so-called Advanced Persistent Threat spectrum. Such code, in combination with weak security practices on network infrastructure, can cause severe damage (see Chapters 16, 17, 20, and 21 in this *Handbook*).

30.6.4 Multiple Security Domains. Many networks implement security solely at the point of entry, where the organization's network connects to the public Internet. Such *monolithic firewalls* are less than effective choices for all but the simplest of small organizations. As a starting point, systems available to the general population should be outside of the internal security domain in a no-man's land between the public Internet and the private internal net such isolation is referred to as a *demilitarized zone* (DMZ). These systems should be afforded a degree of protection by sandwiching them between an outer firewall (protecting the DMZ from the public network) and an inner firewall (protecting the internal network from the public network and controlling the communications between the publicly accessible servers located in the DMZ to the internal network). Alternatively, the DMZ may be a cul-de-sac attached to a separate port on a single firewall, using a different set of traffic management rules. Such a topology permits implementation of the differing security restrictions applicable to the two networks. Systems located within the DMZ should also be suspect, as they are targets for compromise.

Although the industry as a whole agrees on the need for DMZ configurations, it is less appreciated that such restrictions also have a place within organizations. Different groups, departments, and functions within an organization have different security and access requirements. For example, in a financial services firm, the security

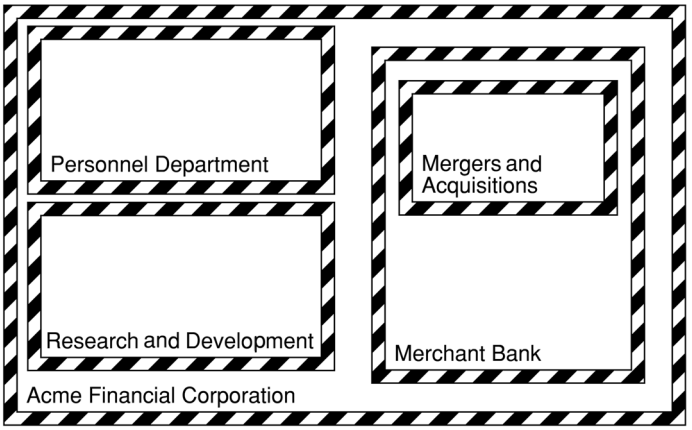


EXHIBIT 30.3 Sibling and Nested Security Domains

requirements differ dramatically among departments. Three obvious examples of departments with different requirements are personnel, mergers and acquisitions, and research and development (see Exhibit 30.3).

The personnel department is the custodian of a wide range of sensitive information about the firm, its employees, and often outsiders who are either regularly on company premises or work regularly with the company on projects. Some of this information, such as residence addresses, pay levels, and license plates, is sensitive for personal or cultural reasons. Other information subjects the organization to legal or regulatory sanctions if it is improperly disclosed or used. In the United States, examples of sensitive data include Social Security numbers, age, sexual orientation, and the presence of human immunodeficiency virus (HIV) or other medical details. Still other information may be subject to specific legal or contractual confidentiality provisions.

The mergers and acquisitions department handles sensitive information of a different sort. Information about business negotiations or future plans is subject to strict confidentiality requirements. Furthermore, the disclosure of such information is subject to a variety of regulations on governmental and securities industry levels. Within the mergers and acquisitions department, access to information often must be on a need-to-know basis, both to protect the deal and to protect the firm from exposure to civil and criminal liability.

Some information in the research and development department is completely open to the public, whereas other information is restricted to differing degrees.

A full implementation of an adequate security environment will require protections that are not only logically different on a departmental basis but also require that different departments be protected from each other. It is difficult, therefore, if not topologically impossible for a single firewall, located at the connection to the outside world, to implement the required security measures. It is difficult to provide assurances that the corporate backbone is free from surreptitious monitoring. In addition, even authorized monitoring can produce communications logs containing sensitive protected information. The network logs then become a hazard in and of themselves.

Securing systems in isolated logical areas is an example of necessary distrust, merely a matter of ensuring that the interactions between the third-party systems and the outside world are allowed to the extent that they are expected. As an example, consider the straightforward situation at Hypothetical Brokerage. Hypothetical Brokerage uses two

trading networks, Omega and Gamma. At first glance, it would seem that that it would be acceptable to place Omega's and Gamma's network gateways on the usual DMZ, together with Hypothetical's Web servers.

However, this grants a high degree of trust to Omega and Gamma and all of their staff, suppliers, and contractors. The most important operative question is whether there is a credible hazard.

Either of the two gateways is well situated to:

- Monitor the communications traffic to and from Hypothetical's Web servers
- Monitor the traffic between Hypothetical and the other, competing network
- Attack the other gateway
- Disrupt communications to and from the other gateway
- Attack Hypothetical's network

Network providers also represent an attractive attack option. A single break-in to a network provider–supplied system component has the effect of compromising large numbers of end user sites. There is ample history of private (PBX) and public (carrier-owned) switches being preferred targets.[42] Provider-supplied SOHO routers with remote management or Wi-Fi capability can also become an attack vector if the provider-supplied configurations are deficient, or if the provider-generated management passwords are compromised.[43]

The solution (see Exhibit 30.4) is to isolate the third-party systems in separate DMZs, with the traffic between each of the DMZs and the rest of the network scrupulously checked as to transmission control protocol and user datagram protocol (TCP/UDP), port number, and source and destination addresses, to ensure that all traffic is authorized. One method is to use a single firewall, with multiple local area network (LAN) ports, each with different filtering rules, to recast Hypothetical's original single DMZ into disjoint, protected, DMZs.

30.6.5 Compartmentalization. Breaking the network into separate security compartments reduces the potential for total network meltdown. Limiting the potential damage of an incident is an important step in resolving the problems.

The same rule applies to the DMZs. For example, in a financial trading or manufacturing enterprise, it is not uncommon to have gateways representing access points to trading and partner networks. Where one places these friendly systems is problematic. Many organizations have chosen to place these systems within their regular DMZ.

Sites have belatedly discovered that such gateways have, on occasion, been found to have acted as routers, taking over that function from the intended routers. In other cases, the gateways have experienced malfunctions and impaired the functioning of the rest of the network (or DMZ). As always, the only solutions certain to work are shutdown or isolation.

Compartmentalization also prevents accidents from cascading. A failure in a single gateway is not likely to propagate throughout the network, because the unexpected traffic will be stopped by the firewall isolating the gateway from the network. Such an event can be made to trigger a firewall's attack alarms.

The network problem is not limited to externally provided nodes. An errant system operating within a noncompartmented network, located as it is within the outer security perimeter, can wreak havoc throughout the entire corporation. Constructing the network as a series of nested and peer-related security domains, each protected by appropriate

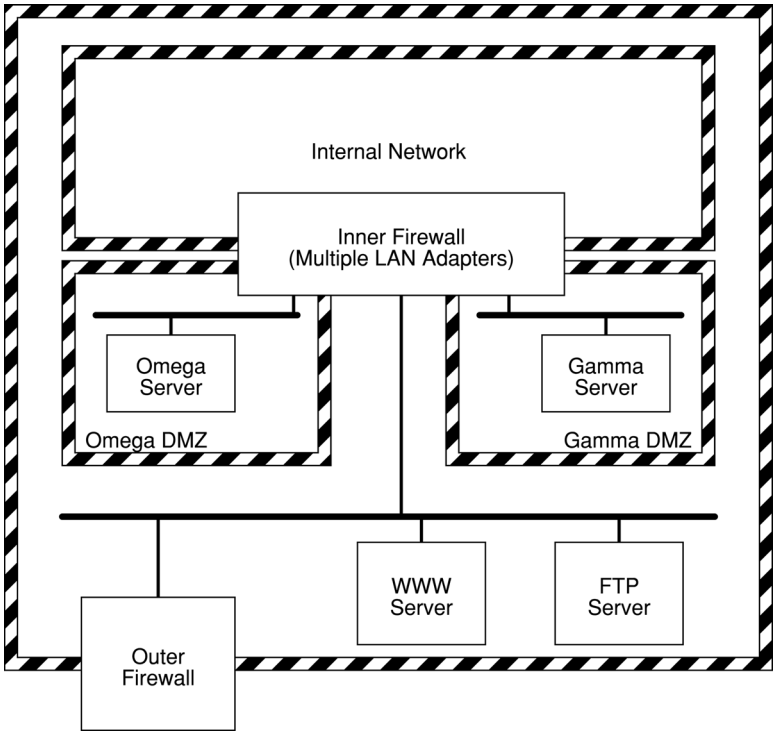


EXHIBIT 30.4 Omega and Gamma Servers in Separate DMZs from Hypothetical's Server

firewalls, localizes the impact of the inevitable incidents, whereas an uncompartmented network permits the contagion to spread unchecked throughout the organization. Larger networks dramatically raise the costs of an incident. It is quite conceivable that the entire budget for compartmenting a corporate network will be less than the cost of a single hour of downtime resulting from the first errant system.

The ready availability of portable storage devices, including USB memory devices and external hard drives, makes compartmentalization an even more serious issue. Exploits have shown the desirability of disabling automatic execution of software stored on plug-in devices.

30.6.6 Need to Access. Sometimes it is easy to determine who requires access, to which resources, and to which information. However, the question of access control often involves painful choices, with many nuances and subtleties. Legal and contractual responsibilities further complicate the question. For example, lack of access may exculpate persons alleged to have misused information.

Physical and logical access controls (see Chapters 23, 28, 29, 32, and 35 in this *Handbook*) need to be implemented for Internet-accessible systems as for any other sensitive and critical system. Controls for such systems must be enforced and respected by all members of the organization. It is important that personnel with restricted access to the network and security infrastructure understand and comprehend the reasons for the security rules and that they not take measures that circumvent those rules. The integrity of an organization's firewalls and network infrastructure is only as good as the

physical and logical security of the personnel, equipment, infrastructure, and systems comprising the firewall and network. Regular auditing of both physical and logical access to infrastructure assets is critical and necessary. Scanning for unauthorized (rogue) Wi-Fi access points is also prudent. As was the past with dial-in modems; rogue Wi-Fi access points can compromise an organization's entire security posture.

The need to maintain information security on communications within organizations argues for the extensive use of security technologies, even when the data packets are never expected to leave the premises. It is sound planning, even within the enterprise, to require applications with privacy or confidentiality requirements to make use of privacy infrastructure such as the Secure Sockets Layer (SSL) for Web-based applications or tunneling such as Layer 2 Tunneling Protocol[44] and Point-to-Point Tunneling Protocol.[45] This helps ensure that sensitive information is limited in distribution.

Needless to say, these approaches should employ high-grade encryption and properly signed X.509 certificates from well-accepted Certification Authorities. Self-signed, expired, and not-generally-accepted certificates should not be used.

30.6.7 Accountability. People often talk about impenetrable systems. Despite the castle-and-moat analogies used in many discussions of security, perimeters are imperfect. Given the likelihood of successful attacks, security personnel must use both technical and managerial measures for effective response.

When securing infrastructure, priority should be given to protective measures that ensure accountability for actions. Just as it is desirable to prevent inappropriate activity, it is even more important to ensure that activities can be accounted for.

For example, there are many ways to carry out denial-of-service attacks. The most troublesome are those which are completely legal. Although some of these attacks, such as distributed denial of service, are belligerent or politically and ideologically motivated, involving remote-control *zombie* programs and botnets (described in Chapter 16), many accidental DoSs can occur without malice in the course of software development. For example, two of the most famous worms that inadvertently led to DoSs, the Morris worm and the WANK worm, were detected inadvertently as a result of implementation errors in their replication mechanisms. These errors caused both worms to proliferate extremely rapidly, effectively producing unintended DoS attacks and subsequent detection.

When a compromised machine is analyzed forensically, the presence of the attack may remain undetected. This is particularly true when the attack involves designer malware not in general circulation (see Chapters 16, 17, and 20 in this *Handbook*).

It is important to analyze security breaches to distinguish among attacks, accidents, and experiments. It is better for weaknesses to be uncovered, even accidentally, within an organization than it is to deal with a truly belligerent security breach. Policies and practices should therefore encourage employees to report accidents rather than try to hide them. As for false alarms caused by overenthusiastic security neophytes, management should avoid punishing those who report illusory breaches or attacks. Accountability provides the raw material to determine what actually happened. The resulting information is the critical underpinning for analysis and education, thus enabling the enterprise to evolve to higher levels of security and integrity.

30.6.8 Read-Only File Security. Many sites allow downloads, usually via FTP and HTTP, of large numbers of files. These files may include copies of forms,

manuals, instructions, maps, and service guides. If such file serving is provided, then it is critical to ensure that:

- The servers supporting the downloading service are secure.
- The contents of the publicly accessible file store are read-only and subject to change control.
- The entire contents of the public file store can be restored quickly in the event of a possible compromise.
- There is a clearly designated party with responsibility for maintaining and protecting the public file service.

30.6.9 Going Offline. Internet-connected systems require responsiveness that is directly related to the out-of-service costs. This may be the cost of lost business; in other organizations, the cost may be that of lost professional time, of damaged public relations, and of lowered morale. In any event, the larger the proportion of the organization (or its customers) affected by the problem, the higher the cost, and the greater the urgency to effect repairs.

Today's interconnected world makes Internet disconnection a truly painful option for a network or security manager. Although the cost of disconnection (or partial disconnection) can run into hundreds of thousands of dollars in lost business and productivity, disconnection in certain situations is both necessary and appropriate. At times, disconnection presents the lowest-cost, most effective way to protect users, systems, and the public. On May 4, 2000, during the epidemic of the "I Love You," Microsoft Outlook–exploiting virus attack, network managers at Ford Motor Company disconnected Ford's network from the outside world to limit the entry of contaminated email into Ford's systems and to prevent Ford's systems from spreading the contagion.[46] The response achieved its goals. Disconnection resulted in less pain than the alternatives. Segmented networks help provide compartmentalization, permitting a more nuanced response than simple disconnection.

The primary issue surrounding disconnection is: What can be disconnected? Who has the authority? Murphy's Law often intervenes. Such important incidents require short response times, inevitably on occasions when senior managers are not available. The best way to provide for this contingency is to furnish guidelines for personnel with authority to defend the systems and a guarantee that actions within the guidelines will be immune from reprisal. If an organization fails to authorize such actions, they also forswear the benefits accruing from the steps to contain damage.

30.6.10 Auditing. In any organization, facilities usage should be monitored and analyzed, including network activity. Because people interested in attacking or subverting the enterprise's networks will attack when they wish, such monitoring and analysis must be part of a continuing process. Ongoing observation and review should include:

- Physical communications infrastructure
- Firewalls, router tables, and filtering rules
- Host security
- File security
- Traffic patterns on backbones, DMZ, and other network segments
- Physical security of systems and communications infrastructure

These areas are both synergistic and independent. They are synergistic in that the combination of all of them is mutually reinforcing in promoting a secure computing environment. They are independent because any one of them may represent a weak point in an otherwise secure environment. For more information on auditing, see Chapter 54 in this *Handbook*.

30.6.11 Emerging Technologies. Emerging technologies continue to recast the challenges in the Internet technology arena. One such challenge is emerging in the use of SSL/TLS-based firewalls.

Using HTTPS as a basis for building encrypted tunnels appears at first impression to be a tremendously enabling technology. Most firewalls permit the initiation of connections from within the protected zone on TCP port 443 for the purpose of allowing connection to the wide variety of Websites that need to secure information, as has been mentioned in other locations throughout this chapter. For privacy and security, such traffic should be opaque to monitoring.

The use of HTTPS as the basis for tunneling also provides a tailor-made technique for abuse. For example, such technologies make it possible for a compromised desktop system to monitor the network, extracting data of interest, and to transmit the resulting data securely to an outside location through the firewall. The opaque nature of the SSL/TLS-based connection makes it impossible to scan the outgoing data.

This technology may well be the death knell of shared-connectivity LANs. It is ever more important to secure the network infrastructure against attacks that allow a rogue machine to accumulate data by monitoring network traffic.

Intrusion detection systems may also need to be repurposed to identify HTTPS connections that have significant traffic volumes, for signs of compromise or abuse. Approaches that treat the entire network as inherently compromised, with the use of IPSec VPN technology from the desktop server, to secure communications infrastructure, may be needed to protect against such attacks and malfeasance. In this context, the use of telnet for console communications is strongly discouraged. SSH and other encrypted protocols should be used.

30.7 ETHICAL AND LEGAL ISSUES. Managing a Website poses a variety of ethical and legal issues, mostly surrounding the information that the site accumulates from processing transactions as well as from performance tracking, problem identification, and auditing. Policies need to be promulgated to ensure that information is not used for unauthorized purposes, and the staff running such systems needs to be aware of, and conform to, the policies. Many of these ethical issues have civil and criminal considerations as well. The topic of information privacy is more completely addressed in Chapter 69 of this *Handbook*. For more information on security policy standards, development and implementation, see Chapters 44, 45, 47, 48, 49, 50, and 66.

30.7.1 Liabilities. The liability environment surrounding Web servers is too new for litigation to have run its course. However, there is no reason to believe that the myriad laws governing the disclosure of personal information will not be fully enforced in the context of Websites.

Websites increasingly handle sensitive information. Financial industry sites routinely handle bank and securities transactions. Email services handle large volumes of consumer traffic, and more and more insurance companies, employee benefits departments, and others are using Websites to deal with extremely sensitive information covered by a variety of regulations.

Part of the rationale for recommending careful attention to the control of Web servers and their supporting systems is the need to create an environment of due diligence, where an organization can show that it took reasonable steps to ensure the integrity, safety, and confidentiality of information.

30.7.2 Customer Monitoring, Privacy, and Disclosure. Customer monitoring is inherently a sensitive subject. The ability to accumulate detailed information about spending patterns, for example, is subject to abuse. A valid use of this information helps to pinpoint sales offerings that a customer would find relevant while eliminating contacts that would be of no interest. Used unethically or even illegally, such information could be used to assemble a dossier that could be the subject of embarrassing disclosure, of insurance refusal, and even of job termination. The overall problem predates the Web. In fact, more than 20 years ago a major network news segment reconstructed someone's life using nothing more than the information contained in their canceled checks, supplemented with publicly available information. The resulting analysis was surprisingly detailed. A similar experiment was reported in 1999, using the Web.[47]

Organizations sometimes violate the most basic security practices for protecting online information, when all of the information required to access an account improperly is contained on the single page of a billing statement. There have been repeated incidents (e.g., CDUniverse, CreditCard.com) where extremely sensitive information has been stored unencrypted on Web-accessible systems. These incidents recur with regularity and are almost always the result of storing large amounts of sensitive client information on systems that are Internet accessible. There is little question that it is inappropriate to store customer credit card, and similarly sensitive data, on exposed systems.

The security and integrity of systems holding customer order information is critical. The disclosure of customer ordering information is a significant privacy hazard. Failure to protect customer banking information (e.g., credit card numbers and expiration dates) can be extremely costly, both in economic terms and in damaged customer relations. Firms processing payment cards are also subject to the Payment Card Industry Security Standards Council's PCI DSS; which can create additional liabilities and obligations on the enterprise.

A Website, by its monitoring of customer activity, will accumulate a collection of sensitive material. It may be presumed that the information is useful only for the site, and is inherently valid, but there are a variety of hazards here, most of which are not obvious:

- Information appearing to originate from a single source may indeed be a compilation of data from multiple sources. Shared computers, firewalls, and proxy servers can give rise to this phenomenon.
- Casual correlations may arise between otherwise unrelated items. For example, it is not an uncommon acceptable business practice for one member of a business group to pay for all expenses of a group of traveling colleagues. Failure to correctly interpret such an event could be misconstrued as proof of illicit or inappropriate behavior.
- Individuals often have a variety of societal roles (e.g., parent, domestic partner, professional, nonprofit involvement). Conclusions drawn from such data may not reflect the individual, but can easily be an unrelated consequence of some other involvement (e.g., an inquiry about the SAT or GRE could indicate an intent to attend a college; it could also mean that a child, niece, neighbor, or other acquaintance could have sought counsel).

The problem with casual associations is the damage they can cause. In the national security area, the use of casual associations to gather intelligence is a useful tool, albeit one that is recognized to have serious limitations. In other situations, it is an extremely dangerous tool, with significant potential to damage individuals and businesses. An example:

> A California-based married businessman flies to New York City. When he arrives, he checks into a major hotel. A short time later, he makes a telephone call, and shortly a young woman goes up to his room and is greeted warmly. Apparently, a compromising situation. The businessman is old enough to be the woman's father. In fact, he *is* her father. That single fact changes apparently inappropriate behavior into a harmless family get-together. Peter Lewis[48] of *The New York Times* correctly notes that this single fact, easily overlooked, dramatically changes the import of the information.

The danger with correlations and customer monitoring is that there is often no control on the expansive use of the conclusions generated. The information often has some degree of validity, but it is both easy to overstep the bounds of validity, and it is difficult, if not impossible, to later correct damage, once damage has been done.

30.7.3 Litigation. The increasing pervasiveness of the Web has led to increasing volumes of related litigation. In this chapter, the emphasis is on litigation or regulatory investigation involving commercial or consumer transactions, and the issues surrounding criminal prosecution for criminal activities involving a Website. More detailed information on this subject appears in Chapter 61 of this *Handbook.*

Civil. Website logs and records can become involved in litigation in many ways. In an increasing number of cases, neither the site owner nor the operator is a party to the action; the records merely document transactions and events involved in a dispute. The dispute may be a general commercial matter, a personnel matter, or even a domestic relations matter involving divorce.

It is important that records handling and retention policies be developed in concert with counsel. The firm's management and counsel also must determine what the policy is to be with regard to subpoenas and related requests. Counsel also will determine what materials are subject to which procedures and regulations. For example, in a case of an email provider (e.g., hotmail.com), material may be subject to the Electronic Communications Privacy Act of 1986 (18 U.S.C.A. § 2510 et seq.). Other material may be subject to different legal or contractual obligations.

Regulatory. A wide range of rules is enforced (and often promulgated) by various regulatory agencies. In the United States, such agencies exist at the federal, state, regional, and local level. (Outside the United States, many nations have agencies only at the national and provincial levels.) Many of these agencies have the authority to request records and conduct various types of investigations. For many organizations, such investigations are significantly more likely than civil or criminal investigations. Regulatory agencies also may impose record-keeping and retention requirements on companies within their jurisdiction.

Criminal. Criminal prosecutions receive more attention than the preceding two categories of investigation yet are much less frequent. Criminal matters are expensive to investigate and prosecute and must pass a higher standard of proof than regulatory or civil prosecutions. Relatively few computer-related incidents reach the stage of a criminal prosecution, although because of its seriousness, the process is the most visible.

Logs, Evidence, and Recording What Happened. The key to dealing effectively with any legal proceeding relating to a Website is the maintenance of accurate, complete records in a secure manner. This is a complex topic, some details of which are covered in Chapter 53 of this *Handbook*.

In the context of protecting a Website, records and logs of activity should be offloaded to external media and preserved for possible later use, as determined by the site's policy and its legal obligations. Once offloaded, these records should be stored using the strict procedures suitable for evidence in a criminal matter. The advent of inexpensive CD-ROM and DVD writers greatly simplifies the physical issues of securely storing such media. A cautionary note about CD-ROM and DVD media is in order: When creating such records, care should be taken to use blank media manufactured to archival standards. Archival quality media is more expensive than the more readily available media of lesser quality; however, archival media has a rated life of decades, cheaper media has a lifespan of a few years.

Archival records not intended for evidentiary use also should be stored offline. This should take the form of physically offline, or at least on systems that are not accessible from the public Internet.

The media should be stored in signed, sealed containers in an inventoried, secure storage facility with controlled access. For this reason, the copies archived for records purposes should not be the copies normally used for system recovery. In the event of an investigation or other problem, these records will be carefully examined for possible modification or misuse.

30.7.4 Application Service Providers. In recent years, it has become increasingly common to outsource entire applications services and hosting. External organizations providing such services are known as applications service providers, more commonly referred to as ASPs. More recently, this has been rechristened as *software as a service* (SaaS). Both raise similar security and integrity concerns. In both cases, significant applications and data are stored outside of the organization, with the organization retaining responsibility for the integrity and confidentiality of these records.

A more bare-boned offering, virtual hosting, has experienced an explosion of popularity, often referred to as *Infrastructure as a Service* (IaaS). Many of the same issues with SaaS apply to environments purchased as IaaS. The dividing line for responsibility is different, but the same classes of problems occur.

Conceptually, ASPs are not new. Many organizations have historically outsourced payroll processing and other applications. Theoretically, the ASP is responsible for the entire application. Often, paying a package price for the application seems attractive: no maintenance charges, no depreciation costs, lower personnel costs, latest technology, and moderately priced upgrades. However, just as a ship's captain retains responsibility for the safety of his ship despite the presence of a harbor pilot, an enterprise must not forget that if something goes wrong, the enterprise, not the ASP, will likely bear the full consequences. In short, the ASP must be required to answer the same questions, and held to the same standards, as an inside IT organization regarding privacy, security, and integrity issues.

The security and integrity issues surrounding the use of ASPs are the same as those surrounding the use of an internal corporate service. Questions of privacy, integrity, and reliability remain relevant, but as with any form of outsourcing, there are additional questions. For example, is stored information commingled with that of other firms, perhaps competitors? Is information stored encrypted? What backup provisions exist?

Is there offsite storage of backups? What connectivity does the ASP have? Where are the ASP's servers, and are they dispersed? What are the personnel practices of the ASP? Does the ASP itself own and operate its facilities, or does it in turn contract out to other providers?

The bottom line is that although outsourcing promises speedy implementation, lower personnel costs, and economies of scale, the customer organization will suffer considerable harm if there is a problem with the ASP, with availability, results, or confidentiality.

In the end analysis, the organization retains liability for its operations and its data. While it is comforting to consider that the legal system will accept "My provider did it" as an excuse for lost, compromised, or untrustable data, it remains an untested theory and is not likely to prevail. Recent experiences with major manufacturers, subcontractors, and tainted products would seem to indicate that outsourcing risk remains a serious potential liability.

Legal process requests, either criminal warrants or civil subpoenas, against ASPs for third-party data represent another hazard. If data belonging to third parties is commingled on the ASP's system, care is required to prevent the unauthorized and inappropriate disclosure of unrelated data belonging to third parties other than the one that is the subject of the request.

None of the items cited are sufficient to justify a negative finding on ASPs as a group. They should, however, serve as reminders that each of the issues related to keeping a Web presence secure, available, and effective apply no less to an ASP than they do to an in-house IT organization.

For information about outsourcing and security, see Chapter 68 in this *Handbook*.

30.8 SUMMARY. Availability is the cornerstone of all Web-related strategies. Throughout this chapter, it has been noted that redundant hosting and routing were necessary to ensure 24/7/365 availability. It was also noted that although some providers of services offer various guarantees, these guarantees almost never provide adequate compensation for consequential damage done to the enterprise. In the end, an enterprise's only protection is to take adequate measures to ensure their own security and integrity.

Operating guidelines and authority are also critical to ensuring the availability of Web resources on a 24/7/365 basis. Systems must be architected and implemented to enhance availability on an overall level. Operating personnel must have the freedom and authority to take actions that they perceive as necessary without fear of reprisal if the procedures do not produce the desired outcome.

Privacy and integrity of information exchanged with the Website is also important. The implementation and operation of the site and its components must be in compliance with the appropriate laws, regulations, and obligations of the site owner, in addition to being in conformance with the expectations of the user community.

Protecting Internet and Web assets is a multifaceted task encompassing many disciplines. It is an area that requires attention at all levels, beginning at the highest levels of business strategy and descending successively to ever more detailed implementation and technology issues.

It is also an area where the smallest detail can have catastrophic impact. Recent events have shown that the lessons of communications history apply. Even a small error in network architecture, key management, or implementation can snowball until it is an issue that can be felt in the corporate boardroom.

30.9 FURTHER READING

Anderson, R. *Security Engineering: A Guide to Building Dependable Distributed Systems*, 2nd ed. Wiley, 2008.

Bhargav, A., and B. V. Kumar. *Secure Java: For Web Application Development*. CRC Press, 2010.

Cobb, Stephen. *Privacy for Business: Web Sites and Email*. Dreva Hill, 2002.

Elliott, E. *Programming JavaScript Applications: Robust Web Architecture with Node, HTML5, and Modern JS Libraries*. O'Reilly Media, 2013.

Gabarro, S. A. *Web Application Design and Implementation: Apache 2, PHP5, MySQL, JavaScript, and Linux/UNIX*, 2nd ed. Wiley, 2013.

Garfinkel, S. *Web Security, Privacy and Commerce*, 2nd ed. O'Reilly Media, 2002.

Habiyaremye, J d'A. *E-Commerce Security Threats*. GRIN Verlag, 2011.

Hassler, V. *Security Fundamentals for E-Commerce*. Artech House, 2000.

Kissoon, T. *Securing Web Applications*. Auerbach, 2013.

Long, F., D. Mohindra, R. C. Seacord, D. F. Sutherland, and D. Svoboda. *The CERT Oracle Secure Coding Standard for Java*. Addison-Wesley Professional, 2011.

McGraw, G., and E. Felten. *Securing Java: Getting Down to Business with Mobile Code*, 2nd edition. Wiley, 1999.

Musciano, C., B. Kennedy, and E. Weyl. *HTML5: The Definitive Guide*, 7th ed. O'Reilly Media, 2014.

Re.think. *The Book on Professional Website Design*. The Bestselling Book.com, 2013.

Scambray, J., Shema, M., and C. Sima. *Hacking Exposed Web Applications*, 2nd ed. McGraw-Hill Osborne Media, 2006.

Sullivan, B. *Web Application Security, A Beginner's Guide*. McGraw-Hill Osborne Media, 2011.

Welling, L., and L. Thomson. *PHP and MySQL Web Development*, 5th ed. Addison-Wesley Professional, 2013.

Zalewski, M. *The Tangled Web: A Guide to Securing Modern Web Applications*. No Starch Press, 2011.

30.10 NOTES

1. J. Swartz, "TJX Data Breach May Involve 94 Million Credit Cards," *USA Today*, October 24, 2007, www.usatoday.com/money/industries/technology/2007-10-24-tjx-security-breach_N.htm
2. J. Vijayan, "TJX Agrees to Pay $40.9M to Visa Card Issuers in Breach Case," *PCWorld*, November 30, 2007, www.pcworld.com/article/140174/article.html
3. T. Dierks and E. Rescoria. "RFC 5246: The Transport Layer Security (TLS) Protocol, v 1.2." *Internet Engineering Task Force, Network Working Group*. Updated by RFCs 5746, 5878, 6176. http://tools.ietf.org/html/rfc5246
4. A. Freier, P. Karlton, P. Kocher. "The Secure Sockets Layer (SSL) Protocol Version 3.0." *Internet Engineering Task Force, Network Working Group*. http://tools.ietf.org/html/rfc6101
5. E. Murray, "Survey on Information Security," *Information Security Magazine* (October 2000).
6. VISA, "PCI Compliance Continued to Grow in 2007," Press Release, January 22, 2008. http://corporate.visa.com/newsroom/press-releases/press753.jsp

7. Dan Kaplan, "Hannaford Tells Regulators How Breach Happened," *SC Magazine*, April 1, 2008, www.scmagazine.com/hannaford-tells-regulators-how-breach-happened/article/108569/
8. Jennifer Thomas, "Cyber Attack Hits Bashas' Chain of Stores," azfamily.com, February 5, 2013, www.azfamily.com/news/Cyber-attack-hits-Bashas-chain-of-stores-189911491.html
9. Ann Bednarz, "Target's Website Woes Getting Fixed, CEO Says," *Network World*, November 18, 2011, www.networkworld.com/news/2011/111811-target-website-253296.html
10. Stephanie Clifford, "Demand at Target for Fashion Line Crashes Web Site," *The New York Times*, September 13, 2011, www.nytimes.com/2011/09/14/business/demand-at-target-for-fashion-line-crashes-web-site.html
11. S. Schiesel and R. Atlas, "By-the-Numbers Operation at Morgan Stanley Finds Its Human Side," *The New York Times*, September 16, 2001, www.nytimes.com/2001/09/16/business/by-the-numbers-operation-at-morgan-stanley-finds-its-human-side.html
12. A. Harmon (2001, October 29) "Breaking Up the Central Office; Staffs Make a Virtue of Necessity," *The New York Times*, October 29, 2001, corrected October 30, 2001. www.nytimes.com/2001/10/29/business/breaking-up-the-central-office-staffs-make-a-virtue-of-necessity.html
13. TRUSTe. www.truste.org
14. P. G. Neumann, "The Green Card Flap," *RISKS Forum Digest* 15, April 18, 1994, http://catless.ncl.ac.uk/Risks/15.76.html#subj2
15. Choe Sang-Hun, "Computer Networks in South Korea Are Paralyzed in Cyberattacks" *The New York Times*, March 20, 2013, www.nytimes.com/2013/03/21/world/asia/south-korea-computer-network-crashes.html
16. John Markoff, "Before the Gunfire, Cyberattacks," *The New York Times*, August 12, 2008, www.nytimes.com/2008/08/13/technology/13cyber.html
17. "Estonia Hit by 'Moscow Cyber War'," *BBC News*, last updated May 17, 2007, http://news.bbc.co.uk/2/hi/europe/6665145.stm
18. G. Keizer, "Hacker Contest Weekend," *Web Design & Technology News*, July 2, 2003, www.webdesignsnow.com/news/070203.html
19. Extensible Markup Language (XML), www.w3.org/XML
20. J. Serjeant, "'Botmaster' Admits Infecting 250,000 Computers," *Reuters*, November 9, 2007, www.reuters.com/article/domesticNews/idUSN0823938120071110; Press Release No. 07-143, United States Attorney's Office, Central District of California, www.justice.gov/usao/cac/Pressroom/pr2007/143.html; FBI, "'Bot Roast II' Nets 8 Individuals," www.fbi.gov/news/pressrel/press-releases/bot-roast-ii-nets-8-individuals
21. Camilla Hall and Javier Blas, "Aramco cyber attack targeted production," *Financial Times*, December 10, 2012, www.ft.com/cms/s/0/5f313ab6-42 da-11e2-a4e4-00144feabdc0.html
22. Shaun Waterman, "South Korea Cyberattack Not Linked to China After All," *The Washington Times*, March 22, 2013, www.washingtontimes.com/news/2013/mar/22/south-korea-cyberattack-not-linked-china-after-all
23. Brian Grow, Keith Epstein, and Chi-Chu Tschang, "The New E-spionage Threat" *Bloomberg BusinessWeek*, April 9, 2008, www.businessweek.com/stories/2008-04-09/the-new-e-spionage-threat

24. Robert Gezelter, "Protecting Web Sites," *Computer Security Handbook*, 4th ed. (2002), Chapter 22, Section 22.3.3, pp. 22–27.
25. TechWise Research, Inc., "Quantifying the Value of Availability," June 2000, http://techwise-research.com/member-resources/white-papers/id/43/quantifying-the-value-of-availability
26. J. Vijayan, "TJX Agrees to Pay $40.9M to Visa Card Issuers in Breach Case," *PCWorld*, November 30, 2007, www.pcworld.com/article/140174/article.html
27. C. Morris, "Sony: PlayStation Breach Involves 70 Million Subscribers," CNBC, April 26, 2011, www.cnbc.com/id/42769019
28. Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address Allocation for Private Internets," IETF RFC 1918, February 1996, http://tools.ietf.org/rfc/rfc1918.txt
29. R. Hinden and B. Haberman, "Unique Local IPv6 Unicast Addresses," RFC 4193, October 2005, http://tools.ietf.org/rfc/rfc4193.txt
30. R. Gezelter, "Security Guidelines for New and Existing OpenVMS Applications," November 12, 1996, www.rlgsc.com/decus/usf96/ad020.pdf
31. N. D. Schwartz and E. Dash, "Thieves Found Citigroup Site an Easy Entry," *The New York Times*, June 14, 2011, page A1, www.nytimes.com/2011/06/14/technology/14security.html?pagewanted=all
32. R. Gezelter, "Internet Dial Tones & Firewalls: One Policy Does Not Fit All," *IEEE Computer Society*, Charleston, SC, June 10, 2003, www.rlgsc.com/ieee/charleston/2003-6/internetdial.html; R. Gezelter, "Safe Computing in the Age of Ubiquitous Connectivity," LISAT 2007, May 2007, www.rlgsc.com/ieee/longisland/2007/ubiquitous.html
33. RFC1597, http://tools.ietf.org/html/rfc1597; superseded by RFC1918 http://tools.ietf.org/html/rfc1918
34. RFC 3927, http://tools.ietf.org/html/rfc3927
35. Robert Gezelter, "Internet Dial Tones & Firewalls."
36. Robert Gezelter, "Safe Computing in the Age of Ubiquitous Connectivity."
37. Formerly RFC 1597.
38. Gezelter, "Internet Dial Tones & Firewalls." Gezelter, "Safe Computing in the Age of Ubiquitous Connectivity." Gezelter, "Internet Dial Tones & Firewalls." Gezelter, "Safe Computing in the Age of Ubiquitous Connectivity."
39. RFC 3927.
40. Robert Gezelter, "Stopping Spoofed Addresses Can Cut Down on DDoS Attacks," *Network World Fusion*, August 14, 2000, www.networkworld.com/columnists/2000/0814gezelter.html
41. Frank da Cruz and Christine M. Gianone, "Using C-KERMIT," 2nd ed. Boston: Digital Press, 1997.
42. M. Slatalla and J. Quittner, "Masters of Deception," New York: Harper Collins, 1995.
43. R Gezelter, "Networks Placed At Risk: By Their Providers," Ruminations—An IT Blog, December 2009, www.rlgsc.com/blog/ruminations/networks-placed-at-risk.html
44. W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn, and B. Palter, "RFC 2661—Layer Two Tunneling Protocol 'L2TP,"' August 1999, http://tools.ietf.org/html/rfc2661

45. K. Hamzeh, G. Pall, W. Verthein, J. Taarud, W. Little, and G. Zorn, "RFC 2637—Point-to-Point Tunneling Protocol," July 1999, http://tools.ietf.org/html/rfc2637

46. K. Timms, telephone interview, summer 2001; K. Bradsher, "With Its Email Infected, Ford Scrambled and Caught Up," *The New York Times*, May 8, 2000, www.nytimes.com/2000/05/08/business/with-its-e-mail-infected-ford-scrambled-and-caught-up.html

47. T. Kelley, "An Expert in Computer Security Finds His Life Is a Wide Open Book," *The New York Times*, December 31, 1999, www.nytimes.com/1999/12/13/business/technology-an-expert-in-computer-security-finds-his-life-is-a-wide-open-book.html

48. P. Lewis, "Forget Big Brother," *The New York Times*, March 19, 1998, www.nytimes.com/1998/03/19/technology/forget-big-brother.html?pagewanted=all&src=pm

CHAPTER 31

WEB MONITORING AND CONTENT FILTERING

Steven Lovaas

31.1 INTRODUCTION 31·1

31.2 SOME TERMINOLOGY 31·2

31.3 MOTIVATION 31·2

31.3.1 Prevention of Dissent 31·3
31.3.2 Protection of Children 31·3
31.3.3 Supporting Organizational Human Resources Policy 31·3
31.3.4 Enforcement of Laws 31·4
31.3.5 National Security Surveillance 31·4

31.4 GENERAL TECHNIQUES 31·4

31.4.1 Matching the Request 31·5
31.4.2 Matching the Host 31·5
31.4.3 Matching the Domain 31·7
31.4.4 Matching the Content 31·7

31.5 IMPLEMENTATION 31·8

31.5.1 Manual "Bad URL" Lists 31·8
31.5.2 Third-Party Block Lists 31·9

31.6 ENFORCEMENT 31·9

31.6.1 Proxies 31·9
31.6.2 Firewalls 31·9
31.6.3 Parental Tools 31·10

31.7 VULNERABILITIES 31·10

31.7.1 Spoofing 31·10
31.7.2 Tunneling 31·11
31.7.3 Encryption 31·11
31.7.4 Anonymity 31·12
31.7.5 Translation Sites 31·12
31.7.6 Caching Services 31·13

31.8 THE FUTURE 31·13

31.9 SUMMARY 31·14

31.10 FURTHER READING 31·14

31.11 NOTES 31·14

31.1 INTRODUCTION. The Internet has been called a cesspool, sometimes in reference to the number of virus-infected and hacker-controlled machines, but more often in reference to the amount of objectionable content available at a click of the mouse. This chapter deals with efforts to monitor and control access to some of this content. Applications that perform this kind of activity are controversial: Privacy and free-speech advocates regularly refer to *censorware*, while the writers of such software tend to use the term *content filtering*. This chapter uses *content filtering*, without meaning to take a side in the argument by so doing. For more on the policy and legal issues surrounding Web monitoring and content filtering, see Chapters 48 and 72 in this *Handbook*.

This chapter briefly discusses the possible motivations leading to the decision to filter content, without debating the legitimacy of these motives. Given the variety of good and bad reasons to monitor and filter Web content, this chapter reviews the various techniques used in filtering, as well as some ways in which monitoring and filtering can be defeated.

31.2 SOME TERMINOLOGY

Proxy—a computer that intercedes in a communication on behalf of a client. The proxy receives the client request and then regenerates the request with the proxy's own address as the source address. Thus, the server only sees identifying information from the proxy, and the client's identity remains hidden (at least from a network addressing point of view). Proxies are widely used by organizations to control the outward flow of traffic (mostly Web traffic) and to protect users from direct connection with potentially damaging Websites.

Anonymizing proxy—a proxy that allows users to hide their Web activity. Typically, such a proxy is located outside organizational boundaries, and often used to get around filtering rules.

Privacy-enhancing technologies (PET)—a class of technologies that helps users keep their network usage private. These include encryption, anonymizing proxies, mixing networks, and onion routing.

Encryption—reversible garbling of text (plaintext) into random-looking data (ciphertext) that cannot be reversed without using secret information (keys). See Chapter 7 in this *Handbook* for details of cryptography.

Mixing networks—users encrypt their outbound messages with their recipient's (or recipients') public key(s) and then also with the public key of a "Mix server." The encrypted message is sent to the Mix server, which acts as a message broker and decrypts the cryptogram to determine the true recipient's address and to forward the message.[1]

Onion routing—an anonymous and secure routing protocol developed by the Naval Research Laboratory Center for High Assurance Computer Systems in the 1990s. Messages are sent through a network (cloud) of *onion routers* that communicate with each other using public key cryptography. Once the temporary circuit has been established, the sender encrypts the outbound message with each of the public keys of all the onion routers in the circuit.

Specifically, the architecture provides for bi-directional communication even though no one but the initiator's proxy server knows anything but previous and next hops in the communication chain. This implies that neither the respondent nor his proxy server nor any external observer need know the identity of the initiator or his proxy server.[2]

A third-generation implementation of onion routing is *Tor* (the onion router).[3]

31.3 MOTIVATION. As a general concept, an individual or group chooses to monitor network activity and filter content through an authority relationship. The most common relationships leading to monitoring and filtering are:

- Governments controlling their citizenry
- Parents and schools protecting their children
- Organizations enforcing their policies

- Governments enforcing their laws
- Governments monitoring terrorists and other nation states

31.3.1 Prevention of Dissent. Probably the most common reason for the opposition to content filtering is that many repressive governments filter content to prevent dissent. If citizens (or subjects) cannot access information that either reflects badly on the government or questions the country's philosophical or religious doctrines, then the citizens probably will not realize the degree to which they are being repressed.

In countries such as the United States, where the Constitution guarantees freedoms of speech and the press, many people feel that only those with something to hide about their activities would fight against censorship. This notion, combined with objectionable content so readily available on the Internet, may be why efforts to question the use of filtering software in libraries and other public places have met with resistance or indifference. Chapter 72 gives many examples of countries using filtering products to limit the rights of their citizens. These examples should be a cautionary tale to readers in countries that are currently more liberal in information policy.

31.3.2 Protection of Children. For the same reasons that convenience stores, at least partially, hide adult magazines from view, the government has required that some classes of information on the Internet be off-limits to children in public schools. This stems from the perception of the schools' role as a surrogate parent and government liability involved in possible failure in that role. A variety of regulations have required the use of content monitoring and filtering in schools and libraries, on the theory that such public or publicly supported terminals should not be using taxpayers' money to provide objectionable content to children. The U.S. Supreme Court has ruled several of these efforts unconstitutional after extreme protest from libraries. Most recently, the Child Internet Protection Act (CIPA)[4] required school districts that receive certain kinds of government funding to use filtering technologies. Most schools have implemented Web filtering, aggressively limiting the Web content available to students. Estimates of the amount of blocked content vary widely, but one school district claims to filter about 10 percent of its total Web traffic due to questionable content.[5]

The obvious target of filtering technology in schools is to keep students from viewing material considered *harmful to minors*. Another less-publicized reason to filter Web content is to prevent students from doing things on the Web that they might not do if they knew that someone was watching. The hope is to keep students from getting a criminal record before they even graduate from high school, protecting them from their bad judgment while they are learning to develop good judgment and learning the rules of society. Arguably, this is the job of a parent and not the job of a school, but schools providing Internet access nonetheless must provide this kind of protection.

High school students and kindergarteners generally have differing levels of maturity, and a reasonable filtering policy would include a flexible, gradual degree of filtering depending on age. Schools could use less-intrusive methods of controlling access to Web content as well, such as direct supervision of students using computers. Parents also have the opportunity to filter the content that their children view at home.

31.3.3 Supporting Organizational Human Resources Policy. Organizations have a variety of reasons for monitoring and filtering the Web content accessed by their employees. The simplest is a desire to keep employees doing work-related activity while at work. Despite studies indicating that the flexibility to conduct limited personal business, such as banking or email, from work produces happier, more

productive workers, managers sometimes view personal use of business computers as stealing time from the company.

A more pragmatic reason to filter some Web content is to prevent "hostile workplace" liability under Equal Employment Opportunity laws. The problem with this approach is that many kinds of content would potentially be offensive to coworkers, so it is difficult to use filtering technology to guarantee that no one will be offended while still allowing reasonable work-related use of the Internet.

Some organizations choose to monitor Web traffic with automated systems, rather than blocking anything, and to notify users of the monitoring. The notification could be in general policy documents or in the form of a pop-up window announcing that the user is about to view content that might violate policy. Either way, the notion is that the organization might avoid liability by having warned the user, but might also avoid privacy and freedom-of-speech complaints. The monitor-and-notify approach also sends the message that the organization trusts its employees, but wants to maintain a positive and productive work environment.

31.3.4 Enforcement of Laws. Law enforcement agencies rarely engage in content filtering, but traffic monitoring is an often-utilized tool for investigating computer crime. Gathering evidence for these investigations often involves catching the traffic as it actually arrives at its destination, so filtering would be counterproductive. In this case, proving identity is the key to getting usable evidence, so privacy-enhancing technologies are real problems. In some cases, the logs of Web proxies—which would identify the real source address of the client machine—are available with a subpoena. Investigations of this sort often target child pornography, drug production, or the theft of computer hardware protected by asset-recovery software.

31.3.5 National Security Surveillance. While government law enforcement agencies work to control activity that violates the nation's laws, other government agencies are concerned with the growth of information and communications technologies as tools of terror and international conflict. From the threat of eventual cyber-war to the more immediate concern with terrorism, governments are exploring ways to detect the use of email, cell phones, social media platforms, and even online games to communicate and coordinate plans for hostile activities. Monitoring these media has become a pressing, expensive, and risky priority for national governments.

As this edition was going to press, secret information was leaked about communications-monitoring activities of the U.S. Government's National Security Agency, rekindling a perennial debate on the balance between privacy and safety in a democratic republic. Despite long-standing precedent and assurances about constitutional protection from unreasonable search, the revelations about the NSA's PRISM and related programs have made it clear that networked communication is fair game in a nation's efforts to stay one step ahead of its adversaries (whether they are real, potential, or even imagined ones).[6] Furthermore, as the U.S. citizenry has discovered, it is very difficult for a government to effectively eavesdrop on hostile foreign electronic communications without gathering *all* communications, at the risk of violating privacy expectations of citizens and allies alike.

31.4 GENERAL TECHNIQUES. Filtering of networked communication can employ two basic tactics: examination of *metadata* and examination of the actual content of a message. The U.S. news media rediscovered the concept of metadata in the wake of the PRISM leak, as NSA representatives insisted that actual message content would

only be examined after filtering out irrelevant communications by looking at phone numbers, call durations, IP addresses, and other transactional information that reveals the source, destination, and temporal context of a conversation.[7] Based on the overwhelming volume of communications traffic traversing the copper wires, fiber-optic cables, and wireless frequencies, it is indeed a daunting task to make filtering (or legal or military) decisions based on examining every word of every message. It us much more efficient, for most purposes, at least to start with these metadata and make assumptions about the relevance or danger of a particular computer, human, or organizational communication partner. Direct examination of the full content of messages is much more resource-intensive, though advances are being made in effective and efficient automated processing of text and images.

31.4.1 Matching the Request. The simplest technique used by filtering technologies is the matching of strings against lists of keywords. Every Web request uses a uniform resource locator (URL) in the general form:

protocol://server.organization.top-level-domain/path-to-file.file-format

Filtering a request for a URL can examine any portion of the string for a match against prohibited strings:

- Filters can match the *protocol* field to enforce policies about the use of encrypted Web traffic (HTTP versus HTTPS). This is more of a general security concern than it is a Web filtering issue, although an organization worried about the need to analyze all traffic could prevent the use of encrypted Web traffic by blocking all HTTPS requests.
- The *server* and *organization* fields describe who is hosting the content. Filtering based in these strings is a broad approach, as it leads to blocking either an entire Web server or an entire organization's Web traffic. See the discussion of server blocking in Section 31.4.2.
- The *top-level-domain* field can be used to filter content; see Section 31.4.3 for attempts to set up an .xxx domain.
- The *path-to-file* field includes the actual title of the requested Web page, so it varies most between individual requests. Whether it is more likely than other fields to contain information useful in filtering depends on the naming convention of the server. This field (and the *file-format* field) is optional, as shown in the request for www.wiley.com, which directs the server to display its default page.
- The *file-format* field tells the Web browser how to handle the text displayed in the page, whether in straight html format or encoded in some other file format (e.g., doc, pdf), or whether to allow dynamic generation of content (e.g., asp). Few filtering products use this field to filter traditional kinds of objectionable content, although enforcing other policies regarding dynamic code in high-security environments can require matching this field.

31.4.2 Matching the Host. Some filtering systems attempt to distinguish between acceptable and unacceptable sources on the Web by inspecting the particular servers or general information portals such as search engines.

31.4.2.1 Block Lists of Servers. Objectionable content tends to be concentrated on individual servers. This naturally leads some organizations to block access to those servers. The two methods of blocking servers are by Internet Protocol (IP) address and by name.

Blocking by IP address simply denies all traffic (or all HTTP traffic) to and from certain addresses. This tactic involves several difficulties based on how addresses are used. First, IP addresses are not permanent. While the numeric addresses of most large commercial servers tend to remain the same over time, many smaller servers have dynamically assigned addresses that may change periodically. Thus, blocking an IP address may prevent access to content hosted on a completely different server than intended. Second, commercial servers often host content for many different customers, and blocking the server as a whole will block all of the contents rather than just the objectionable ones. This is a particular problem for very large service providers like AOL, which grants every user the ability to host a Website. Blocking the AOL servers because of objectionable content would potentially overblock large numbers of personal Web pages. Third, address-based blocking creates the possibility of malicious listing, a practice in which an attacker (or a competitor) spoofs the address of a Web server and provides objectionable content likely to land the server on blocking lists. Some filtering products allow users to submit "bad" sites, providing another opportunity for malicious listing.

Blocking by name involves the Domain Name System (DNS), in which a human-readable name (e.g., www.wiley.com) maps to a computer-readable IP addresses (in this case, 208.215.179.146). DNS allows an organization to change the physical address of its Web server by updating the DNS listing to point to the new IP address, although this does rely on the system as a whole propagating the change in a reasonable amount of time. A device that monitors or filters traffic based on a domain name needs to be able to periodically refresh its list of name-to-address mappings in order to avoid blocking sites whose addresses periodically change. As with address-based blocking, name-based blocking also risks malicious listing as well as both under- and overblocking. Many organizations register multiple names for their servers, for a variety of reasons. For instance, an organization might register its name in the .net, .com, .org, and .biz top-level domains to prevent the kind of misdirection described in Section 31.44 (whitehouse.com). Other reasons for registering a name in multiple domains include preventing competitors from using name registration to steal customers and preventing speculative buying of similar names by individuals hoping to sell them for significant profit. Web hosting companies also provide service to many different organizations, so a variety of URLs would point to the IP address of the same hosting server. Thus, blocking by server name could underblock by not accounting for all the possible registered names pointing to the same server and overblock by matching the server name of a service provider that hosts sites for many different customers who use the provider's server name in the URL of their Web pages.

31.4.2.2 Block/Modify Intermediaries. The Web has become an enormous repository of information, requiring the development of powerful search tools to find information. Early search engines gave way to more sophisticated information portals like Yahoo!, Google, AOL, and MSN. By allowing advanced searches and customized results, these portals give users easy access to information that would be difficult or impossible to find using manual search techniques. Portals have become wildly popular tools for accessing the Internet as a whole; Google claimed 100 million search queries

per day at the end of 2000; in 2013, its site index was estimated to contain over 48 billion pages.[8] Information access at this scale makes portals natural targets for monitoring and filtering. Few commercial organizations prevent their employees from using popular portals, since they have become so much a part of the way people use the Internet. Some countries, however, have blocked access to certain portals for their citizens, hoping to control access to information that might tend to violate national laws (e.g., access to Nazi memorabilia in France) or inspire citizen resistance to the government (e.g., access to information about the Tiananmen Square massacre in China). For further discussion of these issues, see Chapter 72 in this *Handbook*.

31.4.3 Matching the Domain. From 2004 until its final decision in 2011, the Internet Corporation for Assigned Names and Numbers (ICANN) considered and rejected requests for a new top-level domain, .xxx, which would allow providers of sexually related content to voluntarily reregister their sites. Such a domain, it was argued, would be easy to filter in the URL of the request, which presumably would appeal to supporters of Web filtering. It would also allow content providers to show that they were complying with laws preventing children from accessing inappropriate material, by enabling more effective parental filtering. Nevertheless, the move met resistance on both fronts. Conservative religious groups feared that establishing a .xxx domain would legitimize pornography, while not all sexual content providers agreed that the perceived benefit would outweigh either the increased filtering of their sites or the easier monitoring of their clients' traffic.

ICANN rejected several revisions of the .xxx domain proposal over the years, citing the lack of unanimity in the sex content provider community as well as the fear that ICANN might be placed in the position of regulating content, which is outside the organization's charter. Ultimately, however, in 2011, the .xxx domain was approved.[9]

31.4.4 Matching the Content. String matching is simple to do, but difficult to do without both over- and underblocking. For instance, one of the most common categories for content filtering (particularly in the United States) is sex. Blocking all content exactly matching the word "sex" would fail to match the words "sexy" and "sexual." To avoid this kind of underblocking, word lists need to be very long to account for all permutations. A slightly more effective tactic is to block all works containing the string "sex," but this would overblock the words "Essex," "Sussex," and "asexual." Looking for all strings beginning with the combination "sex" would overblock "sexton," "sextet," and "sextant." Simple string matching also ignores context, so blocking "sex" would match in cases where a survey page asked the respondent to identify gender using the word "sex" or in pages describing inherited sex traits or gender roles or sexual discrimination lawsuits.

Other difficulties in string matching involve the vagaries of language. URLs can be displayed in any language whose character set a computer recognizes, so a filter will underblock requests in a language for which it lacks word lists. More generally, in any language it is possible to obfuscate the contents of a site with a seemingly benign URL to avoid filtering. The classic example of this is the pornography site www.whitehouse.com, presumably set up to catch visitors who mistakenly typed "com" when trying to reach the U.S. White House Website (www.whitehouse.gov). More recently, spam marketing campaigns have been setting up Websites linked in email, with meaningless strings of numbers and characters in the URL (e.g., http://2sfh.com/7hioh), making the sites difficult to filter.

In a 2006 study, Veritest compared three of the industry-leading Web filter products (WebSense, SmartFilter, and SurfControl). The winning product underblocked 7 sites, overblocked 8 sites, and miscategorized 10 sites, from a preselected list of 600 URLs. The two competing products fared worse, underblocking 23 and 14 sites, and overblocking 9 and 12 sites out of 600.[10] A meta-study of filtering effectiveness studies performed in 2010 determined that from 2001 through 2008, the average accuracy of filtering products was 78 percent, with some increasing success: During 2007 and 2008, surveys found that the success rate had risen to 83 percent.[11] If this is the performance of the industry's leading edge, then clearly the technology is still developing.

Given the difficulties of accurate matching based on text or address related to the Web page request, a natural alternative is to examine the page content itself. Of course, content matching needs to have access to the unencrypted data in transit, so encrypted Web sessions cause a real problem for this tactic. Some organizations allow (or require) that HTTPS sessions terminate on the organization's own proxy server, potentially allowing the proxy to decrypt the data and perform content analysis.

31.4.4.1 Text. It is possible, although resource intensive, to watch the network traffic stream and look for text that matches a list of undesired content. This sort of matching typically does little analysis of context and so is prone to the same kind of false positives (overblocking) and false negatives (underblocking) described in Section 31.4.2. Moreover, as an increasing amount of Web content involves pictures and sounds, text matching becomes less effective.

31.4.4.2 Graphics. A promising new technique, with applications in visual searching as well as visual content blocking, breaks a graphic image into smaller objects by color or pattern. The technique then evaluates each object against a database of reference images for matching with desired criteria. In the case of blocking objectionable sex content, objects can be evaluated for skin tone and either blocked outright or referred to administrators for manual review if the match is inconclusive. Although content-based filtering has not yet developed into a commercial product, the tools exist and the technology seems applicable not only to still images, but also to video and even audio content.[12] In 2006, the NASA Inspector General's office used an image-search program called Web ContExt to snare an employee who had been trafficking in child pornography.[13]

31.5 IMPLEMENTATION. With the exception of content-based matching, which has not yet reached the market in any significant way, most filtering—whether of address, domain, or keyword—involves matching text lists.

31.5.1 Manual "Bad URL" Lists. Many firewalls provide the capability for administrators to block individual URLs in the firewall configuration. Entered manually, these rules are good for one-time blocking when a security alert or investigation identifies sites hosting viruses or other malware. This approach is also useful for demonstrating the general filtering abilities of the firewall and for testing other Web-blocking technologies. For instance, in an organization using a commercial blocking solution on a Web proxy, a simple URL-blocking rule on the organization's border firewall would provide some easy spot testing of the effectiveness of the commercial solution. Given the extreme size and constant growth of the Web, however, the manual approach does not scale well to protect against all the possible sources of objectionable material.

31.5.2 Third-Party Block Lists. With the enormous size of the Web, the more typical approach is to use a third-party block list. Most of these are commercial products with proprietary databases, developed through a combination of automated "Web crawlers" and human technicians evaluating Websites. Some companies have attempted to prevent researchers from trying to learn about blocking lists and strategies, but the U.S. Copyright Office granted a Digital Millennium Copyright Act (DMCA) exemption in 2003 for fair use by researchers studying these lists.[14] Web-filter companies continue to oppose such exemptions. Two open-source filtering alternatives also exist, with publicly viewable (and customizable) block lists that run on caching proxies: SquidGuard[15] and DansGuardian.[16]

31.6 ENFORCEMENT. Filtering of Web traffic typically occurs either at a network choke point, such as a firewall or Web proxy, or on the individual client machine. Economies of scale lead organizations to filter on a network device, while products designed for parental control of children's Internet use usually reside on individual home computers.

31.6.1 Proxies. A proxy server is a device that accepts a request from a client computer and then redirects that request to its ultimate destination. Proxies serve a variety of purposes for an organization, including reduction of traffic over expensive wide-area network links and Internet connections, increased performance through caching frequently accessed Web pages, and protection of internal users through hiding their actual IP addresses from the destination Web servers. Proxies also represent a natural locus of control for the organization, enabling authentication and tracking of Web requests that go through this single device. Most browsers support manual configuration of a proxy for all Web traffic as well as automatic discovery of proxies running on the organization's network. Organizations that use Web proxies typically allow outbound Web traffic only from the IP address of the proxy, thus forcing all HTTP traffic to use the proxy. Use of an encrypted Web session (HTTPS) is possible through a proxy, although either at the expense of the ability to monitor content (if the proxy merely passes the traffic through) or at the expense of the end-to-end privacy of the encrypted link (if the proxy decrypts and re-encrypts the session).

Individuals also use proxies to maintain the privacy of their activities on the network, as described in Section 31.7.4. Thus, in addition to serving as a natural vehicle for content-filtering applications, proxies also represent a serious threat to those same applications.

31.6.2 Firewalls. A firewall's job is to analyze information about the traffic passing through it and apply policy rules based on that information. Maintaining acceptable response time and throughput requires that the firewall do its job quickly and efficiently. In order to do so, most firewalls merely look at network-layer information, such as source and destination addresses and ports. More recently, firewall vendors have been adding more features to increase security and product appeal. Many companies now call their more advanced firewalls "service gateways" or "security gateways" as the notion of Unified Threat Management (UTM) becomes more popular. These UTM devices combine many features that formerly required individual devices, such as anti-virus, intrusion detection, and filtering of both junk email and Web content.

Sophisticated traffic examination increases the demand on firewall hardware. In order to reduce the performance hit caused by increased packet inspection, many

firewalls allow the administrator to define particular rules or protocols for advanced checking. For instance, since viruses are most prevalent in email, Web, and peer-to-peer connections, the firewall administrator might need to configure only anti-virus checking on rules applying to these protocols. Similarly, if the firewall needs only to monitor outbound HTTP requests from a single IP address, the Web proxy, then the extra processing load of the monitoring function can be constrained to that traffic profile.

The decision between filtering Web traffic at the proxy server (letting the firewall just pass the traffic from that address) and filtering Web traffic at the firewall (having the firewall do the URL inspection) depends on the amount of Web traffic and the budget (one device or two). The decision also affects the strength of the assertion that the organization is successfully filtering objectionable content. If the organization's border firewall performs the filtering, then this assertion depends on the firewall being the only way for traffic to leave the organization's network. Other traffic vectors, including wireless networking, protocol tunneling, and anonymizing proxies, may come into play. If the organization relies on client computers to use a Web proxy by policy, then the degree to which users can circumvent this policy should also be a consideration.

31.6.3 Parental Tools. Although client-based Web filtering is not common in large organizations because of the expense and management of such services on a large scale, products enabling parents to block content for their children at home have become a big business. Many large ISPs, such as AOL and MSN, offer parental content-blocking tools as a free feature of their services. Other companies sell stand-alone products that install on a home computer, with password-protected parental administrative access to content-blocking functions. Net Nanny, CYBERsitter, and CyberPatrol are some of the more popular offerings. These products typically reside on individual computers rather than on a network device, although if a home computer is set up as a hub for network connectivity (such as with Microsoft's Internet Connection Sharing), then the controls can filter traffic in the same way as an organizational proxy server. Many of these products also filter other traffic, including email, peer-to-peer file sharing, and instant messaging, as well as offering foreign language filtering, destination-address blocking, and time-of-day access rules.[17]

31.7 VULNERABILITIES. No security scheme, whether physical or logical, is completely free of vulnerabilities, and Web filtering is certainly no exception to this rule. Users who want to access blocked content have a variety of tactics available to them, although solutions vary in ease of use. IP spoofing, protocol tunneling, and some forms of encryption are not trivial practices, and therefore are the tools of technically adept users in reasonably small number. Other technologies, however, such as anonymizing proxies, translation sites, and caching services, are easy ways for the average user to defeat filtering. Web-filtering vendors are constantly striving to make their products more effective, while privacy and free-speech advocates support ongoing efforts to defeat what they call censorware.

31.7.1 Spoofing. In an organization that performs Web filtering on a proxy server, the organization's border firewall must allow outbound HTTP requests from the IP address of that proxy. A user who can configure traffic to look as though it is coming from the proxy's IP address might be able to get traffic through the firewall without actually going through the proxy. This tactic, known as address spoofing, takes advantage of lax routing policies on routers that forward all unknown traffic to default gateways without checking to see if that traffic came from a direction consistent with

its reported source address. The drawback of spoofing, from the attacker's point of view, is that a large organization with significant amounts of Web traffic traversing the network will notice the temporary unavailability of the proxy caused by the spoofing.

An organization can defeat spoofing by configuring internal routers to check their routing tables, to see if the address of a packet coming into an interface is consistent with the networks available via that interface. The router drops packets with source addresses inconsistent with their actual source. This tactic, called *reverse-path filtering*, requires a more capable (and expensive) router, so it is generally not available for the home user trying to set up network-based protection for parental control.

31.7.2 Tunneling. A more problematic tactic, because it relies on the behavior of applications rather than on subverting network-layer protections, is protocol or application tunneling. An application can encapsulate any other application's information as a packet of generic data and send it across the network. Virtual private network (VPN) clients use this approach to send traffic through an encrypted tunnel.

Protocol tunneling (sometimes called dynamic application tunneling)[18] relies on applications that send data on commonly allowed ports. For example, a user might tunnel a Web session through the secure shell (SSH) application, which uses TCP port 22. SSH often is allowed through firewalls, and because it uses both authentication and encryption, it can be hard to monitor the difference between a legitimate SSH session and a covert tunnel. In this example, the Web browser issues a request for an HTTP page on TCP port 80, and another application running on the client system captures and redirects the port 80 request into an SSH-encrypted tunnel. The other end of the SSH tunnel could be the destination server or a proxy server somewhere between the client and server. The client application (in this case, the Web browser) is unaware of the traffic diversion and needs no altered configuration. This is similar to the approach used by traditional VPNs. A final form of protocol tunneling has come into limited use as IPv6 has become more broadly available. Still in its infancy as an Internet-wide protocol, it nevertheless is available in some pockets of the world, and tunnel brokers offer connections that tunnel IPv6 traffic over the IPv4 Internet backbone. This can become problematic if a home or site believes that it is filtering outbound requests but is only looking for content in IPv4 packets.

Application tunneling (also called static application tunneling)[19] requires reconfiguration of the client application to redirect requests through a different port on the client machine. Typically, the user redefines the destination address for the application to a *localhost* address (in the 127.0.0.x range, referring to the local device), and an application then sets up a connection from that local port to the destination on an allowed network port. This approach requires alteration of the client application's configuration. Some Secure Sockets Layer (SSL) VPN products use this approach to tunnel one or more protocols, or all traffic, across an encrypted HTTP tunnel.

Tunneling via protocol or application generally requires access to either configure existing application settings or install extra software. Thus, tunneling is unavailable to users in organizations that give end users limited control over their computers. Tunneling is more of a problem for parents, whose technically adroit children can install tunneling applications to get around parental controls.

31.7.3 Encryption. Security professionals generally encourage the use of encryption, since it protects sensitive information in transit across a network. However, when users encrypt data to hide transactions that violate organizational policy, encryption can become a liability instead of an asset from the organization's perspective.

In an encrypted Web session using HTTPS (which is HTTP over SSL), the contents of the session are encrypted and thus unavailable for monitoring or filtering based on URL or data content. However, the source and destination IP addresses of the session, which are visible to TCP as the session is set up, remain visible on the network. This is necessary for routers to be able to get the encrypted data packets from one end of the transaction to the other. Thus, while an HTTPS session is immune to filtering by URL text matching or content-based filtering, blocking the destination server is still effective.

As mentioned in the previous section, VPN technologies represent another use of encryption to protect the content of a transaction. Again, although the VPN encrypts the contents of the transaction, the IP addresses of the endpoints must remain visible in order to transport the data to their destination.

Another more complicated version of encryption, called *steganography*, actually embeds data inside other information to avoid detection. For instance, a user could embed a text message within the information used to encode a picture. For more details on steganography and other types of encryption, see Chapter 7 in this *Handbook*.

31.7.4 Anonymity. Most Web monitoring and content filtering relies on the identity of the user. Content filtering can happen without identity information, but unless an organization or country chooses to impose draconian broad-based filtering of all Web requests (and some do choose this approach), the organization might like to enforce filtering requirements for only some users. For instance, a school district might wish to impose stricter filtering on elementary school students than on high school students and use a more permissive policy for teachers and administrative employees.

The bane of this approach is anonymity. When privacy concerns by individuals lead to the use of anonymizing technologies on a scale that makes identity of users difficult to determine, then the only way to comply with policies and laws requiring filtering of content for some groups is to filter for all groups at the level of the strictest requirement. In the case of a school district, if the network cannot distinguish between student traffic and teacher traffic, then the district must impose the requirements of student content-filtering on teachers as well. In fact, many school districts have made this choice, due to both the extreme numbers of hard-to-find anonymizing proxies and the technical difficulty of separating student and teacher use of a common school network infrastructure. Even so, external anonymizing proxies bedevil network administrators' attempts to force compliance with filtering regulations.

Other uses of privacy-enhancing technologies (PET) include network-based anonymity schemes, such as mixing networks and onion routing. The project, known as Tor (originally an acronym for the onion router), has been gaining popularity in recent years. For more information about anonymity and identity, see Chapter 70 in this *Handbook*. For more about PET, see Chapter 42 in this *Handbook*.

31.7.5 Translation Sites. Language-translation sites, such as BabelFish (formerly at http://babelfish.yahoo.com),[20] also offer the possibility of avoiding content filters. The user enters a URL and clicks a button to request a translation of the text of the site. The user's session is between the client computer and Yahoo!, with the requested URL merely passed as keystroke data within the HTTP session, so the potentially blocked site is made available so long as the user can get to Yahoo! This is a special case of a proxy, in that the request typed into BabelFish generates a request to the server at the other end, but the client does not receive the exact results of that reply; the display, instead, includes all the graphics of the original, but the text is translated

into the requested language. So, from the point of view of a monitoring tool, the client is always connected with Yahoo!, and even a content filter looking at text within HTTP packets could be stymied by the foreign language. Recognizing the potential for abuse, Yahoo! published a Terms of Use document for the BabelFish site prohibiting using the service for "items subject to US embargo, hate materials (e.g. Nazi memorabilia), …pornography, prostitution, …[or] gambling items,"[21] among many other classes of activities or products.

Translation sites are also annoying to system administrators because translation of a site in a particular language *into* that language usually results in unchanged content. For example, a translation of an English-language Website from French into English simply passes the content through without alteration.

31.7.6 Caching Services. One of the primary uses for proxy servers has been to reduce network traffic by saving local copies of frequently requested pages. This behavior can circumvent Web filtering, so long as the user can get to the caching server. Google, for example, caches many pages in an effort to provide very fast response to search requests. Large graphic and video files take up the most bandwidth, so Google's Image Search feature often caches them. Often, graphic files are available even for sites that no longer exist. Users wanting to bypass blocks on sexually explicit content often use Google Image Search. Google does provide a feature called Safe Search, with three levels of voluntary filtering. The default (middle) setting filters "explicit images only," while the strict setting filters "both explicit text and explicit images."[22] Google notes, "no filter is 100% accurate, but Safe Search should eliminate most inappropriate material."[23] The Safe Search feature is configurable per user, and offers no password protection, so it is not a Web-filtering technology so much as it is a voluntary sex-based search filter.

31.8 THE FUTURE. As more ways of communicating and distributing content emerge, the content-filtering industry will doubtless evolve to cover the new technologies. At present, vendors sell filtering products for email, Web chat, newsgroups, instant messaging, peer-to-peer file sharing, and FTP, as well as filtering of Web requests. New features will appear in "traditional" Web filtering as well, including filtering of IPv6 tunneled traffic. The latest versions of the home filtering products Net Nanny and McAfee Parental Controls now offer the ability to force safe search options in the major search engines (like Google Safe Search, described in Section 31.7.6), and provide "object recognition," which recognizes certain versions of Web objects (like visit counters) that are commonly used in pornography sites.[24]

Supporters of content filtering, and those who are required to use it and need a reliable product, will be encouraged by the growth of the industry but perhaps disappointed that the problem never seems to be completely solved. Advances in image recognition could provide much better filtering, but may well spawn new ways to alter content to circumvent these tools. Those who decry these products as censorware will point out that, historically, most attempts to censor speech have failed in the end. Eventually, the two sides will probably work out an uneasy compromise, as has happened regarding sales of "adult" content in print and video. As long as some people insist on their right to distribute information that other people find offensive, this conflict is likely to continue. The debate about the balance between freedom and safety will probably become increasingly colored by concerns with illegal activity, terrorism, and cyberwar, which could lead to advances in available tools as well as attempts to define and regulate appropriate use of those tools.

31.9 SUMMARY. For a variety of reasons, some better than others, groups of people with power over, or responsibility for, other groups of people want to control the kind of information to which the other groups has access and to gain knowledge of who is trading information. Monitoring and content-filtering products (or censorware) provide this kind of control using computer technology to examine information flowing across the Internet. Free speech and privacy advocates argue that content filtering prevents legitimate, legal access to information. Even should one grant the legitimacy of filtering in some cases, current technologies are prone to error, both failing to block some objectionable content and blocking some sites that contain no such content.

Most filtering techniques involve examining metadata, matching a string of text or numbers to determine the source or destination of a request or message, or some characteristic of the communication context. This metadata can be used to block access: servers can be blocked by address or by server name. The recent addition of a .xxx domain has opened the possibility of filtering entire top-level domains. Other methods focus on blocking content, examining text or nontextual parts of a Web page, including graphics. Research into image recognition and flesh-tone matching is progressing, and government agencies have used some image-recognition tools in prosecuting cases, but image recognition has not yet entered the commercial market in a large way.

With every protective, or overprotective, strategy comes a group of people dedicated to its defeat. Content filtering has a number of vulnerabilities, chief among which is the use of anonymity via privacy-enhancing technologies such as anonymizing proxies and onion routing. Other ways to defeat Web filtering include the use of protocol and application tunneling, encryption, Web translation sites, and caching services. Filtering technologies have been improving over the years, as has the inventiveness of those dedicated to thwarting them. As information outlets continue to proliferate and new communications media appear, this kind of conflict between protective technologies and privacy-enhancing circumvention is likely to continue.

31.10 FURTHER READING

Brennan Center for Justice, "Internet Filters: a public policy report," 2006. www.fepproject.org/policyreports/filters2.pdf

The Censorware Project, http://sethf.com/freespeech/censorware/essays/censorware org.php

Electronic Privacy Information Center. "Censorware: A Post-CDA Solution?" http://epic.org/free_speech/censorware

Secure Computing. "Best Practices for Monitoring and Filtering Internet Access in the Workplace," www.securecomputing.com/pdf/scc_bestpractices.pdf (note: registration required).

"Seth Finkelstein's Anticensorware Investigations—Censorware Exposed," http://sethf.com/anticensorware

31.11 NOTES

1. D. Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Communications of the ACM* 24, No. 2 (February 1981), www.cs.utexas.edu/~shmat/courses/cs395t_fall04/chaum81.pdf
2. D. M. Goldschlag, M. G. Reed, and P. F. Syverson, "Hiding Routing Information," Workshop on Information Hiding—*Proceedings*. May 1996, Cambridge, UK.
3. www.onion-router.net
4. Codified at 47 U.S.C. §254(h) and 20 U.S.C. §9134.

5. Poudre School District (Fort Collins, Colorado), Information Technology Services, accessed June 1, 2013, www.psdschools.org/department/information-technology
6. "US Intelligence Outlines Checks It Says Validate Surveillance," *The Guardian*, June 16, 2013, www.guardian.co.uk/world/2013/jun/16/nsa-the-nsa-files
7. "NSA Collecting Phone Records of Millions of Verizon Customers Daily," *The Guardian*, June 6, 2013, www.guardian.co.uk/world/2013/jun/06/nsa-phone-records-verizon-court-order
8. "WorldWideWebSize," www.worldwidewebsize.com
9. Internet Corporation for Assigned Names and Numbers (ICANN), "18 March 2011 Draft Rationale for Approving Registry Agreement with ICM's for .xxx sTLD," March 18, 2011, www.icann.org/en/minutes/draft-icm-rationale-18mar11-en.pdf
10. "Websense: Web Filtering Effectiveness Study," January 2006, accessed April 7, 2007, www.lionbridge.com/NR/rdonlyres/websensecontentfilte7fmspvtsryjhojtsecqomzmiriqoefctif.pdf (URL inactive).
11. "Filtering Studies and Their Findings, 2001–2008." Librarian in Black blog, May 7, 2010, http://librarianinblack.net/librarianinblack/2010/05/filtering.html
12. "Using eVe for Content Filtering," eVision Visual Search Technology, accessed June 1, 2013, www.evisionglobal.com/business/cf.html
13. "NASA HQ Raided in Kiddie Porn Probe," *The Smoking Gun*, March 31, 2006, www.thesmokinggun.com/documents/crime/nasa-hq-raided-kiddie-porn-probe
14. S. Finkelstein, "DMCA 1201 Exemption Transcript," April 11, 2003, http://sethf.com/anticensorware/hearing_dc.php
15. SquidGuard Website, accessed June 1, 2013, www.squidguard.org
16. "DansGuardian: True Web Content Filtering for All," accessed June 1, 2013, http://dansguardian.org
17. Top Ten Reviews, "Internet Filter Review 2007," accessed April 4, 2007, http://internet-filter-review.toptenreviews.com
18. SSH Communications Security, "Secure Application Connectivity," accessed June 1, 2013, www.ssh.com/manuals/client-user/53/tunnel-dynamic.html
19. SSH Communications Security, "Secure Application Connectivity."
20. Named after the tiny Babel fish in Douglas Adams's science-fiction classic *Hitchhiker's Guide to the Galaxy* (New York: Del Rey, 1995). The Babel fish, when inserted in a person's ear, would instantly enable the person to understand any spoken language. The fish name is used for a translation technology developed by AltaVista, then owned by Yahoo!, and formerly offered as a free service by Yahoo! It is also the name of a commercial translation firm that has registered the trademark BabelFish.com and which runs the translation service at www.babelfish.com (accessed June 1, 2013).
21. "Yahoo! Search Builder Terms of Use," accessed April 7, 2007, http://help.yahoo.com/help/us/ysearch/ysearch-01.html?fr=bf-home
22. "Google Search Settings," accessed April 7, 2007, http://images.google.com/preferences?q=we+live+together&um=1&hl= en
23. "Google Help Center," accessed April 7, 2007, http://images.google.com/intl/en/help/customize.html#safe
24. Internet Filter Review, "Internet Filter Terms," accessed April 7, 2007, http://internet-filter-review.toptenreviews.com/short-definitions.html

CHAPTER 32

VIRTUAL PRIVATE NETWORKS AND SECURE REMOTE ACCESS

Justin Opatrny and Carl Ness

32.1 INTRODUCTION 32·1
- 32.1.1 Borders Dissolving 32·2
- 32.1.2 Secure Remote Access 32·3
- 32.1.3 VPNs 32·3

32.2 REMOTE ACCESS VPNS 32·4
- 32.2.1 IPsec 32·4
- 32.2.2 Transport Layer Security 32·5
- 32.2.3 User-Authentication Methods 32·6
- 32.2.4 Infrastructure Requirements 32·6
- 32.2.5 Network Access Requirements 32·7

32.3 SITE-TO-SITE VPNs 32·7
- 32.3.1 Multiprotocol Label Switching 32·7
- 32.3.2 Remote-Site VPNs 32·8
- 32.3.3 Information-Assurance Considerations 32·9

32.4 EXTRANETS 32·15
- 32.4.1 Information-Assurance Goals 32·15
- 32.4.2 Extranet Concepts 32·16
- 32.4.3 Types of Extranet Access 32·16
- 32.4.4 Information Assurance Considerations 32·17

32.5 CONCLUDING REMARKS 32·20

32.6 FURTHER READING 32·20

32.7 NOTES 32·20

32.1 INTRODUCTION. The rise of the Internet created a new chapter in human civilization. People are no longer tied to slowly updated information sources such as libraries. The exponential growth in the number of people looking for wide varieties of content also spurred the desire for mobility. If a person can search for information residing half way around the world from home, why not be able to do the same from the local coffee shop or while traveling during a business trip? This information revolution offered an opportunity to provide information and services to consumers, businesses, and employees at virtually any point on the globe and on a multitude of mobile platforms.

The days of focusing on protection at the network perimeter are over and have been for longer than people may realize. Laptops lead the charge of the initial mobile force. Although access to the full processing power of the laptop is useful in some instances, even the lightweight versions are bulky to travel with and have relatively limited battery

life. Smartphones and tablets, fitting in the palm of a hand, a pocket, or a purse, now provide the power to access more commonly used applications.

32.1.1 Borders Dissolving. Widespread Internet access redefined the dynamics of network and perimeter protection. Previously, companies focused on protecting the internal network and systems exposed to the Internet. A perimeter firewall was sufficient to keep the digital predators at bay. The greater challenge is now maintaining the security of the internal network when employees use mobile technologies from home or while traveling. Further complicating the issue is how to allow other business partners—but not intruders—secure access to the systems and information you are responsible for protecting.

Organizations large and small cannot expect to maintain a competitive advantage by restricting employee network and information access to the confines of the workplace. Traveling employees can now maintain direct communications with those at the office. Employees stranded at home due to inclement weather or illness can continue to function, with direct access to necessary information. This mobility provides organizations a viable method to maintain a geographically diverse workforce that allows people to work in closer proximity to customers and partners.

Organizations are also deploying flexible mobile solutions capable of working on smart phones and tablets. Instead of bringing a corporate laptop on a customer visit, a sales associate can use a tablet to present product information, input/change order information, and access other internal resources. There is also a considerable movement to allow employees to use their personal laptop or smart device to access internal resources—and evidence that this policy may actually increase information technology costs instead of decreasing them.[1] This trend increases the potential risk profile, as the organization must provide security at different layers with little to no control over the security protections of the personal device(s) that may connect to internal resources. *Bring Your Own Device* (BYOD) is of increasing importance in discussions of productivity and of security.[2]

As an organization's network of vendors, suppliers, and partners grows, so does the necessity to share information. The organization must look for methods that allow these outside entities to access relevant information without exposing nonrelevant information. Information sharing is not the limit of partner involvement. There may be situations where in-house staff does not have the necessary experience to develop or support certain information systems. Although having a consultant on-staff or on-call may be cost-prohibitive, the organization may opt to allow a vendor to access specific parts of the internal network remotely, to provide the necessary levels of service. These parts of the network are called *extranets.*

Customers continue to demand increasing levels of convenience. The banking industry is a useful example of how consumer demands increase, and how business responds to meet these needs.[3] Instead of having to go to a local bank, the ability to conduct certain inquiries and transactions were possible over the telephone. The limitations of this technology demanded an even greater opportunity, provided by Internet access to accounts and other banking services. Now, the consumer has a visual, point-and-click/touch interface instead of a computerized voice and numeric menus. These same convenience requirements are present in the sale of goods and services. Consumers may not want to go to a brick-and-mortar store to make purchases, but by having an online presence, an organization can meet this demand by providing an e-commerce Website that gives the consumer the ability to view and purchase items at the consumer's time and place of choosing.

32.1.2 Secure Remote Access. Although the implications of not meeting partner and/or consumer demands are obvious, organizations conduct a great deal of planning and review to ensure that meeting these demands does not jeopardize the underlying safety of the information systems that drive the business. The two primary technologies leading to secure remote access are the virtual private network (VPN) and the extranet.

A VPN is a virtual network overlaying an existing set of physical and logical networks. A common VPN implementation is a secured connection allowing a remote client to access an organization's internal network through an *encrypted tunnel*. VPNs also have the ability to extend all or only selected portions of the network to other locations in the organization.

An extranet meets some of the same goals as the VPN with the emphasis being information sharing and e-commerce. Instead of connecting to the internal network, the client establishes an encrypted connection to a Web application server residing in an external service zone. Extranets can introduce additional complexity because the outward-facing systems normally require access to one or more information assets within the internal network.

32.1.3 VPNs. The essence of a VPN is simple: Create an encrypted tunnel into the internal network to protect the transmitted data. However, the concepts underlying VPN technologies are complex and require a great deal of planning to ensure the best possible implementation. VPNs normally fall into two categories: remote access VPNs and site-to-site VPNs. Each VPN type has multiple avenues of implementation, with each having unique requirements.

Allowing people to access internal network resources remotely has significant information-assurance implications. The main information assurance goals of VPNs are securely extending the internal network, protecting data during transmission, and minimizing the security impact of the process.

Smart phones and tablets often also have the ability to establish a VPN tunnel. No matter the type of mobile device, the organization must meet the requirements for remote access while accomplishing the least amount of additional risk.

When properly architected, managed, monitored, and secured, the internal network can provide a relatively safe environment for sharing data. However, the Internet is a hostile environment. The mobile device can no longer count on the internal network protections as a digital comfort zone. Internet hotspots for public access, such as those in restaurants, airports, and hotels, create even greater challenges to information assurance. These public locations (many configured for open access) create sites for malicious actors to intercept network traffic to gather information such as credentials, financial information, and intellectual properly. Properly deployed and managed VPNs and extranets have the ability to thwart this type of monitoring by sending data through encrypted channels.

As the need to increase stricter security measures continues, so does the potential to interfere with the end-user experience. It does not take long for a user to figure out that disabling local security measures can have perceived advantages. Without the client firewall running, the connection speed may increase. Worse yet, by disabling security protections, the user may be able to install and run software otherwise blocked from use. The goal of the security managers must be to maintain high security that is transparent to the user. The process of establishing a VPN tunnel or accessing internal resources should seem easy to the authorized user, with the background security infrastructure

providing the necessary levels of protection without interfering with legitimate users' productivity.

This chapter provides a foundation for understanding of the different types, terminologies, and uses for VPNs and extranets.

32.2 REMOTE ACCESS VPNS. The most common use of a VPN is for secure, client remote access. This allows an authorized user to connect to a VPN and gain access to internal resources from anywhere with an Internet connection, while maintaining the security of the transmission. Remote Access VPNs are waging an on-going battle between traditional Internet Protocol Security (IPsec) and the increasingly popular and powerful Transport Level Security/Secure Sockets Layer (TLS/SSL) VPNs. Each VPN type has distinct advantages and each warrants thorough investigation prior to selecting the technology that best fits organizational needs.

32.2.1 IPsec. IPsec is a suite of network layer protocols designed to set up and protect VPN transmissions. Traditional IPsec VPN implementations use a client-resident application or embedded operating-system service to establish a VPN tunnel into to the internal network.

32.2.1.1 Key Exchange and Management. One of the most complex aspects of IPsec is key exchange and management. IPsec uses Internet Key Exchange (IKE) to facilitate the establishment and management of a Security Association (SA).[4] Key points of IKE Phases 1 and 2 include the following[5]:

- Phase 1 can use main mode or aggressive mode to create the initial IKE Security Association. Main mode—consisting of three pairs of packets—is the most common implementation.
- The first pair negotiates the four-parameter protection suite:
  - Encryption algorithm (e.g., Advanced Encryption Standard, AES),
  - Integrity protection algorithm (e.g., Hashed Message Authentication Mode [HMAC] with Secure Hash Algorithm, SHA-256),[6]
  - Authentication method (e.g., pre-shared key or PKI certificate), and
  - Diffie-Hellman group[7]
- The second pair exchanges encryption keys using Diffie-Hellman.
- The third pair authenticates each side of the connection to the other.
- Aggressive mode accomplishes the same task by only using three packets:
  - The first two messages negotiate the IKE SA parameters and perform a key exchange.
  - The second and third messages authenticate the endpoints to each other.
- Internet Key Exchange Version 2 (IKEv2), defined in RFC 5282,[8] is an updated version, though not interoperable with version 1. This new version attempts to correct some of the shortcomings of the original implementation of IKE.
- Some of the improvements of IKE version 2 include:
  - Substantially fewer RFCs,
  - A reduced number of message exchanges,
  - Fewer cryptographic methods,

- Built-in denial-of-service (DoS) protections,
- State detection and management, and probably most importantly,
- Network Address Translation (NAT)[9] traversal via User Datagram Protocol (UDP).[10]

If an organization largely depends on IPsec, one or more of these improvements may make IKEv2 extremely beneficial and create a much more resilient, higher functioning of VPN service.

- Phase 2 uses quick mode to establish the IPsec source addresses (SAs).
- Each side of the connection will maintain an IPsec SA in its Security Association Database (SAD).
- The initiating device creates and sends its SA proposal to the VPN device.
- The VPN device replies with its SA selection and another hash to authenticate the connection.
- The initiating device then replies with the hash it generates from the previous request.
- If the hash matches the challenge from VPN device, the SA goes into the SAD and the connection proceeds.

32.2.1.2 Authentication Header versus Encapsulating Security Payload. IPsec provides two security protocols for protecting encapsulated data. Authentication Header (AH) protects the integrity of the packet header and payload by using cryptographic hashing to ensure data does not change. Encapsulating Security Payload (ESP) is the more common implementation because it not only provides the integrity protection of AH but also protects the confidentiality by encrypting the entire original packet and creating a new IP header.

AH and ESP can transmit data in either *transport* or *tunnel* modes:

- Transport mode preserves the original IP header information while providing payload integrity and confidentiality payload protection.
- Tunnel mode provides integrity and confidentiality protection for the IP header and payload.

Since transport mode uses the original IP header information, this mode creates incompatibilities with Network Address Translation (NAT) because of layer-four integrity checks. NAT causes the IP address of the packet to change during transit, and in turn, will cause the integrity check to fail. Thus, tunnel mode is the primary method for host-to-gateway and gateway-to-gateway VPN connections.

32.2.2 Transport Layer Security. The Transport Layer Security (TLS) protocol provides a method for protecting client/server communications. One of the most identifiable implementations of TLS is *Secure Sockets Layer* (SSL), which provides the basis for the HTTPS protocol. Although this chapter uses TLS and SSL interchangeably as many products support both, there are some differences. TLS is the next evolution of the Netscape-developed SSL because it is an open-standards protocol supported by the Internet Engineering Task Force (IETF) and provides enhanced security capabilities.[11]

The implementation of TLS/SSL is much less complicated (though not necessarily more secure) than IPsec and provides up to 256-bit encryption capabilities that already exist on virtually all Internet-connected systems. The client provides SSL-related parameters to negotiate an HTTPS connection with the server. The server then replies with the negotiated SSL-related parameters as well as its digital certificate. The client then uses the certificate to authenticate the server. If authentication is successful, the client and server establish the encryption keys used to protect the session and begin encrypted communications.

Although very specific conditions must be present, it is possible to conduct a man-in-the-middle (MITM) attack during the described the negotiation process, rendering the encryption ineffective. See Section 32.3.3.1.7 for more information.

32.2.3 User-Authentication Methods. Before being allowed unfettered access to internal resources, the client must identify and authenticate to the VPN termination device. The simplest authentication method is the password associated with the username as identifier. Some typical methods for validating user credentials are through Remote Authentication Dial-In User Service (RADIUS),[12] Lightweight Directory Access Protocol (LDAP),[13] native Microsoft Active Directory (AD),[14] or Kerberos.[15] To combat the ease of stealing or guessing user names and passwords, most vendors provide alternative or complementary authentication mechanisms. An existing Public Key Infrastructure (PKI)—see Chapters 7 and 37 of this *Handbook*—opens the possibility for using dual-factor authentication by requiring a user or machine PKI certificate. Another multifactor option is to use a cryptographic token such as an RSA SecurID[16] or smartcard.

Access to internal resources should not be wide-open to remote VPN users, nor should a one-size-fits-all mechanism be deployed for VPN users post-authentication. Often applied to remote access VPNs, user authorization based on the extended directory attributes of the authenticated user can further control the level of access to internal assets. For example, based on the user's group membership or role in an LDAP directory, VPN connectivity should be granted only to specific resources on the internal network. Further, a certain class of users or department should be assigned a defined, specific subnet of VPN address space.

Further protections such as downstream firewalls or host-based firewalls may control the remote users' access to resources that should or should not be accessed from offsite by policy. Access control lists (ACLs) applied to the VPN endpoint itself can limit what resources on the internal network may be accessed by remote users as well. VPN users should never be considered of equal security levels; all possible points of control, especially in a role-based access control environment, should be applied.

32.2.4 Infrastructure Requirements. VPN networks are an opportunity to use the External Service Zone. Since the VPN connection device must be Internet facing, zoning requires two different networks connected to the firewall. The external, Internet-facing (VPN untrusted) interface would contain inbound and outbound encrypted traffic only. The second network (VPN trusted) is for unencrypted traffic moving to and from the internal network. By using this principle, the firewall policy would restrict the external, inbound connections to the VPN connection device using only the few required protocols. The firewall policy would restrict all external connections to the unencrypted network while allowing authenticated VPN clients the appropriate level of access into the internal network. This unencrypted network also provides a security inspection point that is unimpeded by encryption.

32.2.5 Network Access Requirements. IPsec and TLS/SSL remote access VPNs provide avenues for secure remote access, but each has unique challenges and opportunities.

32.2.5.1 IPsec. IPsec VPNs typically provide full access to the internal network upon successful connection and authentication. IPSec client VPN implementations using IP protocol 50 (ESP) or IP protocol 51 (AH) do not behave well behind an NAT device, which often limits usefulness for users in hotels, coffee shops, and the like. However, configuring the IPsec to use NAT transversal via UDP or TCP connection may overcome the original limitations.

It is also possible to restrict access to internal hosts and networks through a RADIUS policy, VPN connection device, and/or a firewall policy. One consideration that affects client traffic is *Split Tunneling*.[17] With Split Tunneling enabled, only traffic destined for the internal network flows through the encrypted tunnel. Network traffic bound for the Internet takes the most direct route instead of traveling through the tunnel. This can help to reduce bandwidth by not having nonessential traffic flowing through the tunnel. However, the main disadvantage is losing the ability to inspect the Internet-bound network traffic.

32.2.5.2 TLS/SSL. The original implementation of TLS/SSL VPN was far from its IPsec predecessor. These early VPN connections focused on network pass-throughs to specific hosts and protocols, as well as crude, difficult-to-configure portal of simple links to internal files shares and Websites. The portal concept for remote access is unique to the TLS/SSL VPN. The current portals are more robust and flexible, allowing the administrator to feed specific content and network access depending on a user's classification. The true breakthrough for TLS/SSL VPNs was its match of IPsec's ability to allow full network access, without the necessity for a full client-resident VPN application.

The popularity of these VPNs increased rapidly because they reduce overhead, typically behave well with NATed networks, and are easier to configure than IPsec. Most modern SSL VPN implementations now also support a lightweight remote-access client, some of which are built into current smartphones and tablet devices that replicate the fully managed client/server environment traditionally used in IPsec remote access VPN deployments.

32.3 SITE-TO-SITE VPNs. VPNs are not purely for client remote access. Site-to-Site (S2S) VPNs provide the ability to facilitate internal communications needs. VPN Technologies exist that can create meshed, virtual wide-area networks (WANs) in addition to providing alternatives to traditional WAN implementations.

32.3.1 Multiprotocol Label Switching. Multiprotocol Label Switching (MPLS)[18] continues to extend its dominance in the WAN world. Though it is not a traditional encrypted VPN, it does provide a similar type of service. This section will not cover most MPLS intricacies because deployment must meet organizational requirements, and because each service provider has different offerings. However, MPLS does have several distinct, advantageous characteristics.

Purpose. MPLS creates a meshed, routed virtual WAN network at the service provider level. The MPLS network is then free to route packets directly from one WAN endpoint to another within the confines of its virtual network. WAN sites are able to

communicate directly, increasing the WAN's resilience by eliminating one location from creating as a single point of failure. MPLS networks may also provide multiple levels of quality of service (QoS). QoS provides the ability to prioritize network traffic, allowing certain protocols more bandwidth during high utilization periods. It is also possible to establish IPsec tunnels through the MPLS network to increase the confidentiality of traversing data.

Requirements. As with all WAN routing technologies, the routing hardware and software must support the necessary routing and/or MPLS protocols. In addition, the service provider must already have the technology available. Moving to MPLS may require a significant re-architecture of the existing routing infrastructure.

32.3.2 Remote-Site VPNs. Another common use of an IPsec VPN is for intersite (site-to-site, S2S) communications through the Internet.

Purpose. Although geographic diversity provides the organization closer proximity to customers and materials, traditional WAN implementations may not be available or practical. Remote-site VPNs provide an avenue for bridging this gap.

Alternative WAN. It can be cost-prohibitive to run a traditional WAN connection (e.g., T1/T3, Metro Ethernet, etc.) into a leased or owned location. In addition, even a lower-bandwidth leased line may be overpriced and underpowered to provide direct WAN connectivity to a small branch office. This type of S2S VPN may provide the necessary level of connectivity without the high cost of a leased line. Since this S2S VPN will travel over an Internet connection, it may be possible to provide a higher bandwidth for the price. However, the increased bandwidth may come at the detriment of less predictable network performance such as no end-to-end QoS.

Backup. A S2S VPN over the Internet is also an alternative means for WAN backup. Although a cellular backup card can provide increased bandwidth at a lower cost point than Integrated Services Digital Network (ISDN), the cellular networks typically have higher latency and may be infeasible depending on the building location in relation to a cell tower or provider. For isolated locations, a satellite connection may be feasible but typically has even higher latency than a cellular card. These solutions may be insufficient if the site requires heavy access to systems and information across the WAN. An S2S VPN over the Internet provides a means for higher bandwidth using existing hardware. The S2S connection provides additional redundancy by providing a standby, routable link. If architected properly, the site should notice little to no outage if the primary WAN link fails.

Vendor/Partner Connections. Business needs may necessitate the interconnection of the internal network with a vendor/partner network. It is essential to conduct proper planning, configuration, maintenance, and monitoring of these connections to ensure neither side introduces a negative situation leading to compromise and/or information loss. This setup will typically sit behind a GSD to control access to only necessary protocols, networks, and systems. However, depending on the protocols necessary (e.g., Remote Desktop Protocol, or RDP), it can be difficult to track what occurs from here. The RDP system can act as a jump-off point to other portions of the network not necessarily protected to the same level as they may exist on the internal network.

Requirements. Remote site VPNs require an Internet connection as well as VPN and encryption-enabled endpoints. A typical S2S deployment could be between two routers. Other potential deployments include a router to a VPN-enabled Gateway Security Device (GSD), or from GSD to GSD.

32.3.3 Information-Assurance Considerations. Although the ability for people and remote sites to connect to internal resources provides mobility and increases productivity, the organization must carefully consider the information-assurance caveats. Each implementation will have its own unique requirements, but addressing the main concerns presented in the following sections can provide for more informed decisions.

32.3.3.1 Remote-Access VPN Considerations. Remote-access VPNs provide the most incremental risk. All uncontrolled networks deserve treatment as insecure, thus hostile, computing environments. Since there is uncertainty as to the security of the network the end-user system will be using to connect, it is imperative that the host have sufficient self-protection mechanisms to make up for the loss of the additional defense-in-depth layers provided by the internal network. This understanding will require the security administrator(s) to focus on the necessary precautions to protect all of the organization's systems.

32.3.3.2 Fidelity of the Mobile Device. Since there is no absolute control when mobile systems leave the confines of the internal network, the fidelity of the mobile device is a concern. The risk of compromise at every level—from physical to operating system to applications and beyond—can expose the internal network to increased risk during a VPN session.

There are several ways to reduce the likelihood of a loss of fidelity. Protections may include any combination of host-based security software (e.g., firewall, antimalware, intrusion prevention), current patch levels, and secure configurations. Proper deployment and configuration of these protection measures will help to reduce the possibility of compromise.

Even if the fidelity of the device is acceptable during the initial connection, that does not ensure this status will not change at a later time during the connection. Another method for ensuring the health of the device is through Network Access Control (NAC).[19] NAC provides the ability to interrogate a connecting device before and during the connection to determine such information as patch levels, status of security protections, and memory-resident malware. Unfortunately, NAC is not just a plug-and-play proposition. There are many other implications and requirements—such as determining what to evaluate and/or enforce, inline versus passive deployments and client-based versus client-less solutions—associated with effectively deploying NAC at any point in the network.

32.3.3.3 VPN Client Management. Client administration plays a crucial role in the success of a VPN solution.

IPsec requires the use of a client-side application or embedded operating system mechanism. The primary implications of this include configuring and maintaining the client. The intricacies surrounding a hard-client installation include the necessity of local administrative permissions, potential user interaction, and dealing with a corrupt VPN client application. As with all other applications, updates are necessary due to new enhancements, unsupported older versions, and vulnerabilities. In keeping with the goal of minimizing negative user experience, administrators need a way to update client software with little to no user interaction or side effects. One primary question is to determine whether the VPN vendor has a mechanism to push configuration and other client software updates on connection. Without this functionality, it becomes

more difficult to change client-side parameters. For example, a VPN using a preshared key for VPN authentication gets hard-coded into the VPN software configuration. This makes it difficult to change it without having the end-user enter a new key, in addition to being an avenue for compromise if the key is recoverable.

TLS/SSL VPN clients do not have the same pitfalls as IPsec client management. Instead of a hard client, the TLS/SSL VPN normally uses a small Java or ActiveX-based dynamic client or a lightweight OS-specific installation. In certain VPN systems, the administrator may have the choice to leave the dynamic client resident or remove it on disconnection. If the client becomes corrupted, the user can delete the control, and it will automatically download on the next successful connection.

Having the dynamic client remain resident has two advantages. First, this may provide the ability to have a local, encrypted workspace to store working documents and other files, or to provide access to other resources. Second, this keeps the remote system from having to download the client with each connection. The TLS/SSL client also has an advantage by always checking to see if a resident client is current. If not, the VPN system will force a download of the newest client, helping to minimize client-side management.

Proper monitoring of remote access VPNs is also essential. A compromised VPN connection or credential is extremely dangerous and an increasingly more popular target for attackers. Organizations should monitor logs and connections for suspicious activity and have procedures in place for such problems. For example, if a user ID logs in from the main office in California and less than two hours later the same user ID logs in from halfway around the world, it is likely the second connection is not intended or legitimate. This is especially true of single sign-on (SSO) environments.

32.3.3.4 Protection of the VPN Device. The heart of being able to provide VPN services revolves around the existing network security infrastructure and has direct implications regarding perimeter protection, since the VPN connection devices must be Internet *and* internal facing. These devices must also follow the external service zone discussed in Chapter 26 of this *Handbook*. The device itself is another part of the overall network security profile.

It is important to remove all ancillary exposures. The primary protection method is to configure the firewall to allow access only to the necessary ports on the VPN device. All unnecessary protocols, services, and configuration options should be shut down. If a network management tool uses ICMP and/or SNMP to monitor the VPN device, firewall rules or onboard ACLs should only allow access to these protocols from a specific set of monitoring hosts. Although network reconnaissance and attacks are still possible on these ports, this method drastically reduces the attack profile.

Administrative device access should never use insecure protocols such as Telnet, FTP, or HTTP. Instead, administration should only occur using properly configured encrypted protocols such as Secure Shell (SSH) and Hypertext Transfer Protocol Secure (HTTPS). Administrative access should at a minimum be protected with a user name and strong password, or better yet, add dual factor authentication using a certificate or token combined with a PIN or password. To minimize administrative access, only specific network(s) or host(s) should be able to access the administrative console directly, and if possible, never allowed on any Internet-facing interfaces. To reduce administrative exposures further, it may be possible to use a jump host for all administrative access. However, this creates the potential for a single point of administrative failure if the jump host is not available.

32.3.3.5 Cryptographic Options. Without carefully controlling the VPN's cryptographic protocols, it is possible to choose a combination that creates additional risk of compromise or information disclosure. The strength and proper implementation of the VPN tunnel's cryptographic protection is essential to protecting the data-in-motion and must be sufficient to meet the organization's security requirements. The selection of the available cryptographic protocols for IPsec occurs during the IKE negotiation. The cipher suite determines which cryptographic protocols are available to protect VPN sessions. For example, if 3DES and SHA1 are considered insufficient protection, these and all other weaker alternatives must be removed from the negotiable IKE or cipher suite options. This ensures that incompatible clients are unable to connect and only sufficient protections are in-use during the VPN session.

32.3.3.6 Traffic Inspection. Inspection of network traffic plays a crucial role in keeping the internal network secure. Since the goal of the VPN is to provide secure data transport, this is a hindrance when attempting to evaluate inbound, encrypted traffic. Network traffic inspection can occur on the VPN connection device or on the unencrypted side leading to the internal network. Although the network traffic is one step closer to the internal network, this provides a valid inspection point in which to detect and protect against malicious traffic. Certain VPN devices and GSDs may provide the ability to view data inside the tunnel, evaluate network traffic (providing a troubleshooting point), and inspect data before passing onto the unencrypted network. However, doing all of the inspection on the VPN device can have significant performance implications.

32.3.3.7 Processing Power. Encryption and decryption of VPN tunnels requires sufficient processing power. This issue escalates as the number of clients increases. For situations where there are only a few remote-access clients, it may be possible to use a router to terminate VPN connections and provide security services. The higher the number of remote-access clients, the greater the need for dedicated VPN devices (including GSDs). It is possible to install hardware encryption accelerators to offload cryptographic calculations and management from the shared CPU to a dedicated processing device within the router or GSD, thus easing the processing demands on these units. Larger organizations may have thousands of concurrent users, requiring multiple VPN devices.

32.3.3.8 Interception. Since VPN sessions traverse uncontrolled networks, it is possible for malicious actors to intercept the encrypted communications; however, the datagram protocol underlying TCP/IP makes capturing sequential packets difficult anywhere but the source and destination of the session. If the attackers can capture the data stream, it is increasingly feasible for them to crack the encrypted communications using distributed computing environments with massively parallel processing. However, a more practicable method is an MITM attack while the VPN session is being established. Both IPsec and TLS/SSL VPNs are potentially susceptible to MITM attacks. These MITM attacks typically start with the malicious actor on the same network as the victim and conducting an Address Resolution Protocol (ARP)–poisoning attack[20] to have all traffic flow through the actor's system before being passed on to the rest of the network.

The main IPsec MITM vector is a preshared key (typically a group name and password) to authenticate the IPsec connection itself. This key is typically readily

recoverable from the VPN client configuration file(s), although this approach requires the malicious actor to exploit access to the client side to determine this information. With the group name and password determined, a tool such as *Fiked* will emulate the VPN termination device, authenticate the IPsec connection, and collect the user's credentials as it attempts to authenticate to gain access to the full VPN connection.[21] Although this step is not a compromise of the confidentiality of data-in-motion, it provides the actor the ability to connect to the internal network at will with the harvested credentials. It is possible to defend against this attack using a client or computer certificate for the IPsec authentication.

The ability to compromise the confidentiality of TLS/SSL with relative ease became reality in 2009 with the introduction of *SSLStrip*.[22] SSLStrip uses the ARP-poisoning attack to intercept the clients request to use TLS/SSL. Once intercepted, the tool will establish a valid TLS/SSL sessions with the server and create a rewritten, unencrypted HTTP session with the client. This method does not cause a certificate error message that may alert the user to their being an issue with the connection. The tool also uses a padlock favicon to deceive the user into believing the connection is secure even though the URL is actually using HTTP, not HTTPS. It is possible to mitigate this attack by only allowing direct TLS/SSL connections instead of redirecting from HTTP to HTTPS. In addition, previously established TLS/SSL sessions are not susceptible to this attack.

32.3.3.9 Site-to-Site VPN Considerations

32.3.3.9.1 Infrastructure Design. MPLS can be a wholesale change in how the WAN functions as well as introducing troubleshooting and security implications.

Since an MPLS network can provide any-to-any, mesh functionality, it becomes more difficult to troubleshoot network issues. In a hub-and-spoke WAN topology, it is possible to deploy centrally located probes to monitor and troubleshoot WAN issues. It is unlikely that organizations using MPLS with many locations would invest the time and money to deploy these types of monitors in all locations. However, the service provider may be able to provide additional services that provide this type of troubleshooting capability for an additional fee.

MPLS networks also add to the overall network security burden by requiring an increase in the number of network security devices to provide protection at each individual location. This is largely because the sites can establish direct communications instead of going to a hub first.

MPLS places the security of the WAN infrastructure in the service provider's hands. Ensure the MPLS provider has a robust network and strict processes for keeping MPLS networks separate. One added protection against provider-side errors is to deploy Border Gateway Protocol (BGP)[23] routing security, where both sides of the connection would authenticate before becoming part of the routing table.

Remote site VPNs create additional risk by placing the Internet directly against the internal network at remote sites. If the Internet connects directly to the site's router, there are several protection options. An ACL can restrict communication to this interface only by the other side of the remote site VPN connection. Another option would to be to enable additional security features on the router. However, this may require a code upgrade and increase the processing demands on the router. A more secure but expensive option would be to deploy a GSD that provides security and VPN services—creating an additional another layer of protection between the Internet and the internal network.

Multipoint VPNs are also an important design consideration. Where it has become common, if not necessary, for highly available, high-bandwidth VPN connections, the typical central-site or hub-and-spoke VPN design may not be desirable. Many designs rely on a central point to establish a one-to-many mesh of VPN connections. However, if the central point is unavailable, the entire VPN mesh is down.

Multipoint VPNs remove the dependency on the central VPN endpoint. If the central endpoint becomes temporarily unavailable, the other endpoints are able to communicate and negotiate connections among themselves without the presence of the central (or master) node. This allows the rest of the VPN mesh to stay operational and reroute traffic to other endpoints that would otherwise be unavailable without the central endpoint.

32.3.3.9.2 Cost. MPLS adds cost on many fronts. These include the cost of new or converted circuits and additional services such as QoS or monitoring capabilities. In addition, the equipment running the network must be able to support the new MPLS and routing architecture. There also may be a significant time investment required to redesign the network routing infrastructure.

S2S VPNs using routing devices require sufficient processing power (e.g., random-access memory—RAM—and encryption modules). In certain cases, it may cost more to get to a level of code that supports encryption. The cost of S2S VPN deployments using GSDs increases as the number of security services increases, as well as the number of clients that can connect. If deploying GSDs to remote sites, there is not only a monetary cost, but also higher administrative costs for managing this new piece of the routing and security infrastructure.

32.3.3.9.3 Availability. VPNs are less of a convenience and more of a necessity. A VPN outage can cause severe interruptions for a highly mobile workforce. Even with layers, it can be difficult to eliminate all points of failure due to the associated costs. If the organization considers VPN as a necessity, high-availability solutions are a requirement.

One possible solution is to load-balance inbound connections to distribute them across multiple devices. If one device fails, the other devices will continue to service previously connected clients and will provide a place for new clients to connect. Another option is active/standby failover: If one device fails, the other will be available for new connections. A preferable method would be for the active member to replicate connection information to the standby member. If the active member fails, the standby member can continue servicing the existing VPN tunnels and can accept new connections. This whole strategy also hinges on the main VPN connection points having redundant Internet links, uninterruptable electric power, environmental controls, and back-end network infrastructure.

Cloud-based VPN is also becoming a popular alternative to provide primary or backup availability of VPN remote access.[24] These services eliminated the need for a dedicated VPN infrastructure to be built and maintained by organization. Instead, a network may be extended logically to a cloud provider that integrates these connections into a distributed, highly scalable VPN infrastructure hosted in the cloud. This class of service is popular when there are many geographically diverse VPN endpoints or remote users, where low latency is necessary or uptime is critical.

32.3.3.10 Implications of Elusive VPNs. VPNs and VPN-like services occur in more places than most people realize. These unseen VPNs range from legitimate

VPN services such as GotoMyPC[25] to malicious VPNs such as botnet command-and-control channels and reverse tunnels. Although VPN sites such as GoToMyPC can be useful, it is up to the organization to decide if this is an acceptable means of *ad hoc* remote access. Some of the key issues to consider with this type of VPN service are that it bypasses internal security controls such as packet and content inspection, and the third party is creating and managing the VPN connection.

Malicious VPNs are a source of continuing concern. Botnet command and control channels can use encryption to evade network defenses that protect information. A reverse tunnel works by redirecting network traffic destined for a listening port to the same port on another host. For example, if a compromised FTP server had an active reverse tunnel, any user attempting to connect to that FTP server would redirect to another FTP server. Beyond the compromised host theory, malicious people already on the internal network may establish an outgoing VPN connection in an effort to hide network traffic.

The most elusive of all are applications that provide embedded VPN-like services. Peer-to-Peer (P2P) networks tunnel file sharing through proprietary and encrypted protocols. Skype[26] not only has the ability to create encrypted voice calls, but also has VPN-like capabilities that allow encrypted file transfers. Network defenses should be able to detect the setup of these sessions and stop them before the connection is completed.

It is up to the organization to develop internal network security defenses and policies necessary to protect against all elusive VPNs.

32.3.3.11 Impact of IPv6. Mass migration to IPv6 is still occurring at a slower pace than one may expect, even though IPv6 was developed in the late 1990s and the allocation of the last remaining available IPv4 address blocks to regional Internet registries occurred in 2011. As of the writing of this *Handbook* in early 2013, only about 12 percent of networks in global BGP tables are advertising IPv6 prefixes. Although most organizations continue to ignore the risks and ramifications of delaying enterprise IPv6 transition plans, it is likely that the protocol is already in their environment. Almost all modern operating systems have dual IPv4 and IPv6 network stacks, usually enabled by default, and typically prefer IPv6 by default. IPv6 and its associated transition technologies have specific implications for VPNs.

Without ubiquitous end-to-end IPv6 connectivity, there are several IPv6 transition technologies (such as *6to4*[27] and *Teredo*[28]) that allow IPv6 capable systems at both ends to tunnel communication over legacy IPv4 networks. The 6to4 method encapsulates the IPv6 packet in the payload of IPv4 protocol type 41 packets and requires the use of relay routers along with at least one globally assigned IPv4 address.[29] Teredo encapsulates IPv6 packets in the payload of UDP datagrams and also uses relay routers. UDP allows for better traversal of environments using NAT.

These technologies lessen the barrier to adoption of IPv6 by enhancing reachability of IPv6-enabled services; however, these technologies are also often not well understood or managed. Most network and security administrators learn the hard way that when working as designed, a connection (VPN or local) to an IPv6-enabled resource over a transition technology may traverse the organization's border more than once—adding severe latency and passing data through several relay routers not controlled by the organization.

These transition technologies have serious security impacts on both the local network and VPN connection—misconfigurations and lack of experience being the most

dangerous. Without properly configuring (and testing) both endpoints, it is possible for the IPv6 traffic to take an unsecure path before (or even after) entering the VPN tunnel, putting data in the clear and therefore at risk. It is critical for network administrators and VPN users to have a valid, tested IPv6 configuration to ensure IPv6 traffic (both native and tunneled) is traversing *only* the secured VPN tunnel and not traversing any encapsulated tunnel not controlled by the organization. Unless this threat is managed, data may be vulnerable even when an organization thinks its VPN is providing added security to a remote site or worker.

These transition technologies, many of which enabled by default, also tend to set up tunnels and connections automatically. Endpoint configuration, especially on client workstations, must be centrally configured or managed to disable these automatic connections or to specify specific configurations so that unintended tunnels are not configured.

IPv6 also has native support for IPsec (both AH and ESP) through the use of extension headers[30] (versus proprietary IPsec clients), though its use is optional. Configuration and use of IPsec in an IPv6 environment requires the same discipline in choosing and configuring cryptographic options as IPv4. It is also possible to use IPv4 IPsec to protect IPv6 transition technology tunnels, as unencrypted tunnel sessions would otherwise be vulnerable to interception and/or manipulation without this protection. Both of these options are generally misunderstood and often unconfigured on networking and VPN devices. Organizations should only consider native IPsec for IPv6 after a thorough education and planning process.

32.4 EXTRANETS.

As the necessity for convenient access to data increases, so do the challenges associated with providing the protections required to secure this type of access.

32.4.1 Information-Assurance Goals.

Allowing external entities to access internal data remotely has major information-assurance implications. The main information-assurance goals of extranets are protecting shared information assets, preventing information exposure, and minimizing ancillary risks.

Protecting Shared Information Assets. Much of the current information security market focuses on protecting hosts and networks from digital threats. Although the loss of a server or portion of the network undoubtedly causes issues, there may be areas of greater concern. Organizations develop and sustain competitive advantage by manipulating data collected or created into valuable and actionable information. By mitigating the risks associated with external access, there is less chance for a breach of information security to occur.

Preventing Information Exposure. The concept of least privilege is essential when designing system access. One would not expect a financial analyst to have access to proprietary product specifications. The same concepts hold true when assigning access to external entities. Identity management provides an accountability mechanism for providing rights and monitoring actions of issued accounts. Access management provides the enforcement mechanism that keeps those accounts from viewing or changing data not explicitly necessary. The fusion of these two principles reduces the risk of inappropriate or unexpected access.

Minimize Ancillary Risks. It does not take long for Internet-facing systems to come under attack. Minimizing Internet attack vectors occurs from the Network Layer to the Application Layer. The extranet servers are not all-purpose systems. The network protection policies should restrict access to these services so that the Internet

community only has access to the specific service(s) advertised. Extranet services require greater consideration, as network security devices still lag in protecting the protocols and customizations routinely occurring at the Application Layer.

32.4.2 Extranet Concepts. ***External Service Zone***. As described in Chapter 26 of this *Handbook*, extranet networks are another opportunity to use zoning, largely because extranets require robust protections and make extensive use of resources in other zones. Zoning allows administrators to be more granular on the services allowed into the network, as well as providing another location for additional security mechanisms.

N-Tier Architecture. If an extranet system provides all aspects of a service, there exists a single point of failure and compromise. N-Tier architecture provides the ability to distribute different aspects of the service offering to additional systems. The extranet system will provide the front-end user interface; the application server facilitates the logic and processing necessary to access data; and the back-end database server stores and provides the data. Since the extranet server resides on an external service zone, the firewall policy must allow that server to connect to its next tier servers on the internal network. When possible, restrict access from the extranet server to the internal network to only necessary systems and services. This architecture helps to create a more robust infrastructure that reduces unnecessary exposures.

TLS/SSL Encryption. The most visible marker of TLS/SSL encryption is the little golden padlock favicon that appears in an Internet browser when connecting to a server using HTTPS. When a host attempts to connect to an extranet server, several steps occur before the encrypted session can begin. The connecting host provides SSL-related parameters to negotiate a HTTPS connection with the server. The server then replies with the negotiated SSL-related parameters and a digital certificate. The client then uses the certificate to authenticate the server. If authentication is successful, the client and server establish the encryption keys used to protect the session and begin encrypted communications.

32.4.3 Types of Extranet Access. The type of systems and services able to go onto the extranet continues to grow. Providing extranet services is always subject to the needs of the organization, and the following sections provide some common examples of extranet systems.

Vendor/Partner Information Sharing. Businesses survive and thrive on competitive advantage, and increasingly need to share information with strategic partners. By providing access to internal information, these partners can better understand and meet the organizational needs. Enterprise Resource Planning (ERP) systems contain the data that serves as the digital lifeblood of an organization. On the extranet, vendors and business partners could have access to work with information stored in the ERP system.

E-Commerce. Conducting business transactions digitally can create efficiencies and reduce costs. E-commerce occurs at the business and consumer level. Businesses can use Electronic Data Interchange (EDI), which allows two entities to exchange standardized, digital versions of documents such as a purchase order and payments without human interaction. The extranet provides a means to allow business partners to exchange EDI data without granting access to the internal servers that actually process and store data. Consumers access e-commerce Websites to view and purchase goods and services.

Employee Productivity. Providing extranet services is not restricted to business partners and consumers. It is possible to enhance employees' access to specific internal services and systems by allowing access without requiring a full VPN connection. Email is the universal business tool, and providing Webmail services can allow employees to access email from almost anywhere in the world on any device with an Internet browser. Employees may find it helpful to be able to access and change benefits information from home. This may also be an opportunity to allow access to intranet information through a content management system.

Outsourced. Enterprises not possessing the requisite internal human or technological expertise to achieve the organization's IT goals have long looked to outsourced solutions to meet their needs. Software as a service (SAAS) provides offerings such as productivity applications, collaboration, and email (e.g., Microsoft Office365) and CRM (e.g., salesforce.com). Infrastructure and platform as a service (IAAS and PAAS, respectively) provide on-demand storage and computing (Amazon Web Services—S3/EC2 or Rackspace Open Cloud). Each of these outsourced offerings continues to mature and redefine how enterprises develop, manage, and present their information.

32.4.4 Information Assurance Considerations

32.4.4.1 Technical Security. Securing an extranet is not possible at a single layer or point in the network. Each layer is interdependent and requires specific protections to ensure the security of the extranet network.

32.4.4.2 Infrastructure. Since the extranet sits behind a GSD, network layer security measures are normally the protections encountered first.

Traffic Inspection. As reiterated in multiple parts of this and other chapters, some security devices have a distinct inability to evaluate encrypted network traffic. Extranets can create additional issues as the encrypted session may travel from the client directly to the extranet server. This potentially requires traffic inspection to occur on the extranet server. Although possible, this method will increase the processing requirements of the server.

Another option is to terminate the SSL connections on an upstream device. This device would then use unencrypted protocols on the downstream side when communicating with the actual extranet server. This method adds an additional layer of complexity when troubleshooting, but it can relieve the extranet server of the encryption and security inspection processing burdens.

Internal Network Exposure. Although the compromise of an extranet server is significant, it should not provide unfettered access to internal resources. Although the firewall policy allows access to specific extranet services from the Internet, the firewall policy must also granularly restrict access from the extranet server to internal systems.

32.4.4.3 Server. Server-level protections are the true first line of defense. It is difficult for network-layer security mechanisms to protect a server that is advertising a vulnerable service. System hardening is an approach to reducing the overall risk profile of the server. Nonessential protocols and services are shutdown. The operating system (OS) may have additional parameters such as file and account permissions that further reduce the risk profile. A script can help to ensure these changes occur in a consistent manner. One well-known system-hardening script is Bastille Linux.

Due to the complexities of operating systems and applications, vulnerabilities will occur. The simplest method of remediating vulnerabilities is to update the server with a patch. It is important to test patches for adverse effects before deployment to the production environment. This practice will help to reduce the potential for unexpected results and downtime.

Vendors can take weeks and months to develop and release a patch. In certain cases, third-party patches will become available for unpatched, actively exploited vulnerabilities. Although the ability to use these patches is an option, it is inadvisable due to the inability to verify their integrity and safety.

Even though the goal is to minimize the risk profile of the server through all of these means, it may also be necessary to provide additional server-level protections, including host-based firewall and intrusion-prevention software. These protections have the advantage of being contextually aware of operating systems and application functions, unlike those residing at the network layer.

Virtualization is already a deeply entrenched aspect of internal server infrastructures and on the extranet. Although virtualization provides economies of scale by allowing multiple extranet systems to run on a single server or cluster, there are many security implications. Without clustering, a single point of failure exists if the physical server goes down. Since the virtual switch directs most of the layer 2 traffic generated by the VMs, there is minimal native visibility and security enforcement of the network traffic generated between the VMs. VM-embedded security enforcement is continuing to become more readily available to protect at the hypervisor level. There are also ongoing concerns of escaping virtual sessions and allowing unauthorized access to another virtual session.

32.4.4.4 Application. Some common application layer vulnerabilities include buffer overflows, Structured Query Language (SQL) injection, and cross-site scripting (XSS). Buffer overflows existed decades ago but continue to be a favorite amongst malicious code developers. This condition exists because the application does not do proper bounds checking on input fields, or uses functions that fail to do the same. SQL injection is a malicious attempt to insert an SQL query into server-side requests. By failing to check for and protect against this type of attack, the back-end server may return the information the malicious query requests. XSS is a type of code injection commonly used for phishing. A malicious actor may send an unsuspecting person a link that redirects some or all content elsewhere.

Vulnerabilities at the application layer can easily defeat even the most secure network layer configurations. Applications developed by multibillion-dollar organizations can be just as vulnerable as something developed internally in a small facility. Developers must understand and follow secure coding practices to help minimize the number and severity of application vulnerabilities. See Chapters 38, 39, and 52 in this *Handbook* for discussions of programming and application security, and Chapters 30 and 31 for Web-based security.

32.4.4.5 Policies. Since extranets provide external entities access to internal information and systems, it is advisable to have a policy that governs the use of these systems. The policy may include information about requirements such as needing a confidentiality agreement before allowing a business partner access to the extranet or mandating the use of TLS/SSL encryption for all extranet communications. Policies do not provide an active protection mechanism but do establish expectations when using extranet systems and consequences of not following those requirements.

32.4.4.6 Access and Identity Management. Access management is crucial to ensuring that only necessary and relevant information gets to the requesting user. The permissions given to this user will dictate the data the user is able to access and should follow the concept of least privilege. Identity management ensures that only properly credentialed entities can gain access to the extranet and underlying systems.

Typically, an external user will authenticate using a username and password. Unfortunately, this is a poor way to validate an entity's identity because of the ease of stealing or guessing passwords. A more reliable method for identifying an external entity would be to require multifactor authentication using a username/password combination and issuing the user a digital certificate. The certificate would then become a part of the connection and authentication process. Although not impossible, it is vastly more difficult to spoof and/or steal a digital certificate.

Federated Identity Management (FIdM)[31] is a common method to provide single sign-on authentication capabilities between two or more partner organizations. Protocols such as Security Assertion Markup Language (SAML)[32] create a standards-based method for each side to communicate and validate authentication information. Due to the intricate details of FIdM, this chapter only mentions its potential use as an authentication mechanism for extranet services. Since each federated organization must place a certain level of trust in the other, it is important to go through a thorough planning process to identify and mitigate the risks associated with this technology.

Chapters 28 and 29 of this *Handbook* contain detailed information about identification and authentication.

32.4.4.7 Availability. Extranet systems are an important part of business operations. When extranet access is unavailable, certain transactions do not occur that can cause issues at points throughout the organization. Extranet availability is a function of the network and server infrastructures.

A single Internet connection has an implied single point of failure if that connection goes down. If the network infrastructure leading to the extranet only has a single path to follow, this becomes another single point of failure. By adding redundancy to the network infrastructure, such as a second Internet connection and secondary routing paths, there is less likelihood of downtime.

Server infrastructures also have the same potential to be a single point of failure. Instead, using a single server to provide an extranet service, it is possible to deploy a cluster of servers and load balancing. The cluster will protect the availability of the service if one or more servers go down. In addition, this provides the ability to perform server maintenance with little to no effect on availability. Load balancing helps to optimize the number of connections going to each cluster member.

32.4.4.8 Impact of IPv6. The largest impact of IPv6 on the extranet is infrastructure and application support. As noted previously, operating systems and applications must be able to accept, process, and reply to the IPv6 packets; the routing infrastructure must be able to communicate using IPv6 natively or through an IPv6 transition tunneling mechanism (e.g., Teredo or 6to4); and lastly, the security infrastructure must be able to evaluate IPv6 packets. A breakdown in any one of these will render the extranet vulnerable or inaccessible.

32.5 CONCLUDING REMARKS. VPNs and extranets have many powerful features that can enhance an organizations ability to conduct and improve employee mobility and business functionality. Support and input must come from all levels of the organization to ensure that the secure remote access systems meet the organization's specific needs and expectations. Most important, these systems must adhere to and enhance, not diminish, the organization's overall security profile.

32.6 FURTHER READING

Stewart, J. Michael. *Network Security, Firewalls, and VPNs*. Jones & Bartlett Learning, 2010.

Whitman, Michael E., Herbert J. Mattord, and Andrew Green. *Guide to Firewalls and VPNs*, 3rd ed. Delmar Cengage Learning, 2011.

32.7 NOTES

1. Tom Kaneshige, "BYOD Planning and Costs: Everything You Need to Know," *CIO*, December 13, 2012, www.cio.com/article/723864/BYOD_Planning_and_Costs_Everything_You_Need_to_Know
2. Caroline Baldwin, "BYOD Increases Productivity, but IT Departments Need To Be Prepared." *ComputerWeekly*, August 2, 2012, www.computerweekly.com/news/2240160757/BYOD-increases-productivity-but-IT-departments-need-to-be-prepared
3. Nasim Z. Hosein, "Internet Banking: Understanding Consumer Adoption Rates among Community Banks," Academic and Business Research Institute Website, January 25, 2010, www.aabri.com/LV2010Manuscripts/LV10038.pdf
4. C. Kaufman, P. Hoffman, Y. Nir, and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2): Request for Comments 5996," Internet Engineering Task Force, September 2010, http://tools.ietf.org/html/rfc5996
5. S. Frankel, K. Kent, R. Lewkowski, A. Orebaugh, R. W. Ritchey, and S. R. Sharma, "Guide to IPSEC VPNs: SP 800-77," NIST Special Publications, December 2005, http://csrc.nist.gov/publications/nistpubs/800-77/sp800-77.pdf
6. S. Kelly and S. Frankel, "Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec: Request for Comments 4868," Internet Engineering Task Force Datatracker, May 2007, http://datatracker.ietf.org/doc/rfc4868
7. M. Friedl, N. Provos, and W. Simpson, "Diffie-Hellman Group Exchange for the Secure Shell (SSH) Transport Layer Protocol: Request for Comments 4419," Internet Engineering Task Force, March 2006, http://tools.ietf.org/html/rfc4419
8. D. Black and D. McGrew, "Using Authenticated Encryption Algorithms with the Encrypted Payload of the Internet Key Exchange Version 2 (IKEv2) Protocol: Request for Comments 5282," Internet Engineering Task Force, August 2008, http://tools.ietf.org/html/rfc5282
9. Internet Engineering Task Force, "Network Address Translators (NAT)—Description of Working Group," Internet Engineering Task Force Datatracker, November 2001, http://datatracker.ietf.org/wg/nat/charter

10. J. Postel, "User Datagram Protocol: Request for Comments 768," Internet Engineering Task Force, August 28, 1980, https://tools.ietf.org/html/rfc768
11. T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol v1.2—Request for Comment 5246," Internet Engineering Task Force, August 2008, http://tools.ietf.org/html/rfc5246
12. A. DeKok and G. Weber, "RADIUS Design Guidelines: Request for Comments 6158," Internet Engineering Task Force, March 2011, http://tools.ietf.org/html/rfc6158
13. K. Zeilenga, "Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map (Request for Comments 4510)," Internet Engineering Task Force, June 2006, https://tools.ietf.org/html/rfc4510
14. Microsoft, "So What Is Active Directory? (Windows)," Windows Dev Center—Desktop, 2013, http://msdn.microsoft.com/en-us/library/windows/desktop/aa746492(v = vs.85).aspx
15. Jason Garman, *Kerberos: The Definitive Guide*, O'Reilly Media, 2003.
16. EMC^2, "RSA SecurID," EMC^2 Security, 2012, www.emc.com/security/rsa-securid.htm
17. Tom Shinder, "Split Tunneling versus Force Tunneling for DirectAccess Clients," Microsoft TechNet, March 30, 2010, http://social.technet.microsoft.com/wiki/contents/articles/split-tunneling-versus-force-tunneling-for-directaccess-clients.aspx
18. Internet Engineering Task Force, "Multiprotocol Label Switching (mpls): Charter for Working Group," Internet Engineering Task Force Datatracker, November 2012, updated July 2013, http://datatracker.ietf.org/wg/mpls/charter
19. J. Sabir, "Network Access Control (NAC) Vendor List," JafSec.Com Website, 2013, http://jafsec.com/Network-Access-Control/Network-Access-Control-A-B.html
20. Corey Nachreiner, "Anatomy of an ARP Poisoning Attack," WatchGuard Infocenter, 2011, www.watchguard.com/infocenter/editorial/135324.asp
21. Irongeek.com, "FIKED," Irongeek.com Website, August 4, 2009, www.irongeek.com/i.php?page = backtrack-r1-man-pages/fiked
22. Irongeek.com,"Using SSLStrip To Proxy an SSL Connection and Sniff It," Irongeek.com Website, 2013, www.irongeek.com/i.php?page = videos/sslstrip
23. CISCO, "Border Gateway Protocol," CISCO DocWiki, October 16, 2012, http://docwiki.cisco.com/wiki/Border_Gateway_Protocol
24. VPN Reviews, "Best VPN Reviews—Compare and Find VPN Account Providers," Best VPN Reviews Website, 2013, www.vpnreviews.com
25. GoToMyPC, "Access Your Mac or PC from Anywhere," Citrix GoToMyPC Website, 2013, www.gotomypc.com/remote_access/remote_access
26. Skype, "Wherever You Are, Wherever They Are—Skype Keeps You Together," Skype, 2013, http://beta.skype.com/en
27. B. Carpenter and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds: Request for Comments 3056," Internet Engineering Task Force, February 2001, www.ietf.org/rfc/rfc3056.txt
28. Microsoft, "Teredo Overview," Microsoft Technet Windows, January 15, 2007, http://technet.microsoft.com/en-us/library/bb457011.aspx
29. Carpenter and Moore, "Connection of IPv6 Domains."

30. Charles M. Kozierok, "IPv6 Datagram Extension Headers," The TCP/IP Guide Website, September 20, 2005, www.tcpipguide.com/free/t_IPv6DatagramExtensionHeaders.htm
31. EDUCAUSE, "Seven Things You Should Know about Federated Identity Management," EDUCAUSE, September 2009, http://net.educause.edu/ir/library/pdf/EST0903.pdf
32. Thomas Hardjono, Nathan Klingenstein, Anil Saldhana, Hal Lockhart, and Scott Cantor, "OASIS Security Services (SAML) TC," *OASIS: Advancing Open Standards for the Information Society*, 2013, https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security

CHAPTER 33

802.11 WIRELESS LAN SECURITY

Gary L. Tagg, CISSP and Jason Sinchak, CISSP

33.1 INTRODUCTION 33·2

33.1.1 Scope 33·2
33.1.2 Corporate use of Wireless LANs 33·2
33.1.3 Functional Benefits of Wireless 33·3
33.1.4 Security Benefits of Wireless 33·3
33.1.5 Centralized Management 33·4
33.1.6 Overview and History of the IEEE 802.11 Standards 33·4

33.2 802.11 SECURITY FUNDAMENTALS 33·6

33.2.1 Terminology 33·6
33.2.2 Authentication and Access Control 33·7
33.2.3 Data Confidentiality 33·8
33.2.4 Key Management 33·8

33.3 IEEE 802.11 ROBUST SECURITY NETWORK 33·9

33.3.1 Features 33·9
33.3.2 802.1X Overview 33·10
33.3.3 EAP, EAPoL, and PEAP 33·11
33.3.4 Insecure Legacy EAP protocols 33·12
33.3.5 Detailed EAP/EAPoL Procedure 33·12
33.3.6 RSN Key Hierarchy and Management 33·16
33.3.7 Temporal Key Integrity Protocol (TKIP) 33·17
33.3.8 Counter-Mode/CBC-MAC Protocol (CCMP) 33·19
33.3.9 Fast Secure Roaming 33·20

33.4 FUNDAMENTAL WIRELESS THREATS 33·21

33.4.1 Unauthorized Network Extensions 33·21
33.4.2 Layer 1 Attacks 33·23
33.4.3 Endpoint Attacks 33·24

33.5 SPECIFIC WIRELESS SECURITY ATTACKS 33·24

33.5.1 MAC Address and IP Spoofing 33·24
33.5.2 EAP (WEP/WPA/WPA-Enterprise) RADIUS Impersonation Attack 33·25
33.5.3 Management Frames 33·26
33.5.4 Hotspot and Cafe Latte Attacks 33·27
33.5.5 WEP 33·29
33.5.6 WPA/WPA2 Pre-Shared Key 33·31
33.5.7 WPA/WPA2: Wi-Fi Protected Setup (WPS) 33·32

33.5.8 MS-CHAP-v2—Divide and Conquer 33·33

33.6 MITIGATING CONTROLS 33·34

33.6.1 Improvements 33·34
33.6.2 Authentication Server and Client Certificates 33·35
33.6.3 Endpoint Supplicant Configurations 33·36
33.6.4 Wireless Intrusion-Detection Systems 33·37

33.7 SECURE ENTERPRISE DESIGN 33·40

33.7.1 Benefits of Wireless Controller Architecture 33·41
33.7.2 Network Segmentation 33·42
33.7.3 User Segmentation 33·44

33.8 SECURITY AUDITING TOOLS 33·44

33.9 CONCLUDING REMARKS 33·46

33.10 ABBREVIATIONS AND DEFINITIONS 33·46

33.11 FURTHER READING 33·50

33.12 NOTES 33·50

33.1 INTRODUCTION. IEEE 802.11 wireless local area networks (LANs) are now ubiquitous and have major benefits such as mobility, flexibility, rapid deployment, and cost reduction over traditional wired networks. However, as with any networking technology, wireless LANs create opportunities for unauthorized individuals to potentially access the enterprise network and the information carried over it.

This chapter provides an overview of wireless LAN technologies, security threats and attacks, and how to address them. It is structured as follows:

- 802.11 history and technological overview
- 802.11 security fundamentals
- Detailed coverage of 802.11 security covering both the original legacy functionality and the upgraded system first defined in 802.11i-2004
- Fundamental wireless-medium security threats
- Specific technical wireless-LAN security attacks
- Technical mitigating controls for specific security attacks
- Overarching secure enterprise design principals

33.1.1 Scope. The scope of this chapter is the security of ANSI/IEEE standard 802.11 wireless LANs. This chapter does not consider any other wireless systems, such as mobile telephone networks, or other wireless standards such as HomeRF, Bluetooth, WiMax, or HiperLAN. This chapter provides a high-level overview of fundamentals with technical deep dives in specific areas necessary to comprehend threats. Many of the basic terms and concepts are documented in Chapter 32 in this *Handbook*.

33.1.2 Corporate use of Wireless LANs. Corporations have been using wireless LANs since the 1990s. However, to begin with, the market was fairly small and the technologies proprietary. In the late 1990s and early 2000s, the groundwork was laid for the mass adoption of wireless LANs. The starting point was the publication of ANSI/IEEE standard 802.11 that provided a baseline design enabling manufacturers to develop interoperable products at lower costs.[1]

33.1.3 Functional Benefits of Wireless. The main advantages of implementing wireless networks are mobility, flexibility, and cost reduction.

- **Mobility:** Wireless technologies enable staff to access network information via mobile terminals as they move around the office campus; examples include warehouses, shop floors, and hospitals. Within an office environment, wireless technologies provide a flexible alternative or addition to the wired network. Often, desks and meeting rooms have a limited number of Ethernet connections; wireless technologies can cost effectively provide additional network connections as required.
- **Flexibility:** Public wireless networks (hotspots) allow staff to utilize idle time between meetings, in airports, coffee shops, and even on airplanes in flight.[2] Typical uses include access to the corporate LAN, along with information on the Internet. Public hotspots can also be trunked over the enterprise wireless LAN to provide internet access for consultants and visitors.
- **Cost Reductions:** Costs can be lowered by not having to install physical network links between buildings separated by a road, river, railway tracks, or even a city block. A wireless link can be set up between two buildings provided there is an uninterrupted line of sight between them. In addition, economies of scale can be realized with a wireless medium. A single access point (AP) (thick or thin) can service one or many end users and scale appropriately by bumping excessive users to neighboring APs. This capability is exclusive to wireless. A wired network has a fixed capacity for per-port access and virtual LAN (VLAN) assignment. The use of wireless with Service Set Identifier (SSID) VLAN support can reduce the networking hardware volume. In the case of dynamic VLAN assignment within a wired LAN, endpoint or server moves require proper VLAN pruning at the switch layer to retain security yet provide the required VLANs. Cost savings can be realized within a wireless LAN by reducing network-management tasks and overall complexity to maximize resources and hardware usage.

33.1.4 Security Benefits of Wireless

- **Physical Security:** An AP and supporting hardware can be hidden from end users to protect it from physical attack. A wireless AP can be hidden above the ceiling in contrast to physical endpoint network jacks, which must be accessible by all users who need access to the internal network. This leaves the door open for unauthorized access to a network port in the absence of enterprise wide 802.1X port security.
- **Segmentation Visibility:** Wired networks commonly assign VLANs on a per-port basis, or in advanced configurations, using one per Media Access Control (MAC) address. In port-based VLAN assignment, management relies heavily on proper access-layer switch configuration and the pruning of sensitive VLANs not necessary for a specific business unit or location. MAC address-based VLAN assignment requires the use of MAC authentication-aware switch hardware and a backend authentication and RADIUS server for MAC → VLAN mappings. MAC-based VLAN assignments can improve management capability and oversight, but pose a significant security risk as they are susceptible to MAC-address spoofing.

In a wireless environment, VLANs can be assigned on a per-SSID basis and administrative efforts can be greatly reduced. A user is no longer confined to a particular

physical location for access to a necessary VLAN, as long as the SSID is available on accessible APs. In the event of expansion, deploying a new AP in close proximately will deliver the necessary VLANs with the corresponding SSID that a client supplicant profile is configured to search for. Since VLAN assignment can be controlled through minor client supplicant configurations and backend RADIUS privileges, segmentation can be determined during the asset-provisioning process and require relatively minor management thereafter.

In an environment which relies heavily on access-layer switches for endpoint connectivity, it is not uncommon to find switching or upstream router protocols making their way to access-layer switch ports. These include protocols such as Dynamic Trunking Protocol (DTP), VLAN Trunking Protocol (VTP), Hot Standby Router Protocol (HSRP), Cisco Discovery Protocol (CDP), Open Shortest Path First (OSPF) routing, and so on. A user with port access can exploit multiple known vulnerabilities in these protocols.

33.1.5 Centralized Management. Thin-client AP environments leverage wireless controllers to provide a central configuration for all associated APs. Thin-client APs are configured on the controller, and intelligence is provided by a backend user directory (Extensible Authentication Protocol–Remote Authentication Dial In User Service, or EAP-RADIUS) and controller itself. The widespread yet central connection of a mesh of APs can additionally be leveraged for security monitoring of the wireless airwaves in a Wireless IDS fashion as described in Section 33.4.

33.1.6 Overview and History of the IEEE 802.11 Standards. The first IEEE 802.11 standard was published in 1999. Following publication, work continued on developing 802.11 with amendments published on a regular basis. The publication of the 802.11b standard increased WLAN bandwidth from 2Mb/s to 11Mb/s, making it a possible technical replacement for a wired LAN. Following that were the 802.11a and 802.11g standards that increased throughput further to 54Mb/s.

These and other amendments were then brought together in the revised standard 802.11-2007. Development continued via further amendments, which were once again brought together with the issue of 802.11-2012.[3]

The first version of the standard contained authentication and confidentiality services in an attempt to provide similar levels of security as wired LANs. The authentication service consisted of two systems called *Open Authentication and Shared Key Authentication*, and the data confidentiality service was called *Wired Equivalent Privacy* (WEP). These services are often referred to as the *legacy security services*.

In 2004 the 802.11i standard was released and defined the *Robust Security Network* (RSN) system, which provided enhanced authentication and confidentiality services. The authentication service has two options: the first is the use of pre-shared keys targeted for homes and Small Office Home Office (SOHO) users and the second uses the 802.1X/EAP framework for enterprise use.

Following the issue of the 802.11-2007, amendments were introduced to provide security mechanisms for mesh wireless networks and the use of 802.11 for transportation applications. There were also enhancements to the core 802.11i algorithms to secure some management frames, and enable fast transition to other APs for time critical services such as Voice over IP (VoIP). Throughput was further improved to a maximum of 600Mb/s by the release of 802.11n-2009.

Exhibit 33.1 provides an overview of the development of the 802.11 standard, the amendments, and identifies those with security functionality using shading.

EXHIBIT 33.1 802.11 Standards

Standard	Description	Security Functionality Description
802.11-1999	The original ANSI/IEEE standard 802.11, 1999 edition. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.	Defined open and shared key authentication and WEP
802.11a-1999	An extension to the 802.11 standard that provides an improved throughput of up to 54Mb/S in the 5GHz frequency band.	None
802.11b-1999	The ANSI/IEEE 802.11b specification is an extension to the 802.11 standard that specifies a higher speed physical layer (up to 11Mb/s).	None
802.11d-2001	802.11d extends the original 802.11 standard to define physical layers, for adoption in additional countries.	None
802.11e-2005	802.11e extends the original 802.11 MAC to provide quality of service support for time-critical applications.	None
802.11g-2003	This standard provides an improved throughput of up to 54Mb/S in the 2.4GHz band.	None
802.11h-2003	802.11h provided spectrum and transmit-power enhancements to the 802.11a 5GHz standard to allow its use in Europe.	None
802.11i-2004	802.11i enhanced the 802.11 MAC to provide improvements in security. This standard defined the RSN security system. The two main encryption algorithms are also known as WPA and WPA2.	Defined the RSN security system
802.11j-2004	802.11j enhanced the 802.11 standard, to add channel selection for 4.9 GHz and 5 GHz in Japan.	None
802.11-2007	The ANSI/IEEE standard 802.11-2007 rolled up all the above amendments into a revised standard.	Security functionality defined in 802.11i-2004
802.11k-2008	802.11k enhanced 802.11-2007 to improve the management and performance of wireless networks. Access Points collect information from clients to control channel selection, roaming, and transmit power control.	None
802.11n-2009	802.11n provides an improved throughput of up to 600Mb/S in both the 2.4 and 5GHz bands.	None
802.11p-2010	802.11p enhanced the 802.11 standard for intelligent transportation system applications.	Dedicated security functionality defined in the standard

(*continued*)

EXHIBIT 33.1 (*Continued*)

Standard	Description	Security Functionality Description
802.11r-2008	802.11r provided fast transitions between APs and to establish the required Quality of Service before the transition.	Redefined the security key negotiation protocol to avoid the lengthy 802.1X process
802.11s-2011	802.11s defines a self-configuring multihop wireless distribution system to provide Mesh networking.	Defines a password-based authentication and key establishment protocol known as SAE
802.11u-2011	802.11u provided improvements to the user experience with wireless hotspots, addressing issues such as enrollment, network selection, and authorization for roaming.	None
802.11v-2011	802.11v extends the management of attached wireless clients established in 802.11k to include the ability to configure the clients.	None
802.11w-2009	802.11w introduced protection for some wireless management and control frames to address potential denial of service deauthentication spoofing attacks through the use of a Message Integrity Code (MIC).	Extended 802.11i to also protect some management frames
802.11y-2008	802.11y produced an enhancement to the existing 802.11 standard to enable high-powered wireless LAN operations in the 3650-3700 MHz band for the USA.	None
802.11z-2010	802.11z defines Tunneled Direct Link setup that enables client devices to establish a direct link with each other whilst remaining associated to an AP. This can more than double throughput between the clients.	None
802.11-2012	The ANSI/IEEE standard 802.11-2012 rolled up all the above amendments to 802.11-2007 into a revised standard.	RSNA defined in 802.11i and mechanisms designed in the above amendments

33.2 802.11 SECURITY FUNDAMENTALS. This section describes the underlying 802.11 security fundamentals in preparation for the details covered in later sections.

33.2.1 Terminology. To better understanding this chapter, readers will find clarification of a few commonly misunderstood 802.11 terms. Section 0 has a more extensive glossary.[4]

Authentication: Authentication is the first step in a two-step process of client connection to an AP. This step validates the client's authority to start associating

with the AP. Network-based authentication such as username and password takes place after layer 2 authentication and association.

Association: Association is the process of the AP accepting the client connection and allocating resources for it. This includes things such as adding client specific information such as supported data rate, data protocol 802.11 b/g/n, and MAC address information.

802.1X: 802.1X provides the encapsulation of Extensible Authentication Protocol (EAP) implementations within 802 communication mediums. 802.1X does not define a specific authentication method, but provides a vehicle for EAP implementations and their underlying methods. 802.1X and ultimately 802.11i are major components of the modern 802.11 RSN system. 802.1X allows or denies client access to requested resources until the client is successfully authenticated.

RADIUS: *Remote Authentication Dial In User Service* (RADIUS) is a network protocol used for authentication, authorization, and accounting (AAA). Historically, RADIUS servers leveraged a flat-file directory for user-based access decisions, but modern implementations leverage a dedicated directory such as Windows Active Directory, Lightweight Directory Access Protocol (LDAP), or a rational database such as Microsoft SQL Server or Oracle. In a Robust Security Network (RSN), a RADIUS server is only responsible for brokering the authentication/authorization of a user requesting access to an AP. Authentication data is transparently sent from the client to the RADIUS server by the AP, and the RADIUS server leverages an external directory to determine a response.

SSID versus BSSID: An SSID is a unique identifier used by a client to establish connectivity to a particular wireless network. An AP can provide multiple SSIDs on the same channel through the use of the same or multiple interfaces. A BSSID is the unique identifier for a Basic Service Set (BSS). A BSS consists of an AP and associated clients or clients.

An **Extended Service Set** (ESS) is a series of BSSIDs (AP interfaces) sharing the same SSID. This helps enable a wireless client to seamlessly move between APs using the same SSID. A BSSID is a separate interface with its own MAC address; multiple SSIDs can share the same interface and MAC address. On a commercial AP the first SSID/VLAN pair will use the BSSID interface/MAC address and each SSID after that will use a virtual MAC address, which increments the BSSID by a small value for each SSID. A BSSID can be used to reference a unique interface or AP (assuming the AP only has one interface). Depending on the vendor of a particular AP, a single BSSID will send broadcast beacons with all SSIDs in one sweep. If an SSID is placed on its own BSSID, it will have a dedicated beacon, which may improve compatibility.

33.2.2 Authentication and Access Control.

In the absence of 802.1X port security, wired-LAN access control is primarily reliant on physical security. To access the LAN, an attacker first needs to have physical access to a connection point. In contrast, the nature of wireless LANs means any wireless client within radio range can potentially connect to the internal LAN network, bypassing physical controls.

To address this issue 802.11 provides an optional native-authentication service. The legacy 802.11 standard includes two protocols, open authentication and shared key authentication. Neither of these protocols was adequate for secure access.

Open authentication is a null-authentication service allowing any and all clients to connect and associate. Shared key authentication (SKA) requires the client to use a cryptographic key to successfully authenticate. This method did put a lock of sorts on the network, but it was soon exploited by attackers who acquired keying material to help compromise the system.[5] SKA lacked unique authenticated user tracking, and forced an out-of-band key management process. SKA has since been deprecated and should not be used, except for necessary backwards compatibility with older devices. The legacy 802.11 security system did not provide any access-control functionality. Although most equipment included access filtering based upon MAC address, this control was not part of the standard and the legacy service is therefore easily defeated.

The RSN defined by 802.11i is a much stronger system, and provides multiple mechanisms to authenticate the device and the user.[6] The system defines a personal profile for home/SOHO use based on PSKs, as well as an enterprise profile based upon the 802.1X/EAP framework, which allows the use of a backend authentication server. The 802.1X authentication system also allows for access control decisions to be taken by the network to restrict network resources available to an authenticated user through technologies such as VLAN segmentation.

33.2.3 Data Confidentiality. Data confidentiality on a wired network is provided by physical security and layer-2 boundaries, which limit the accessibility of the data. Unless an attacker is physically connected to the wired LAN and logically resides between a sender and recipient on the network, through ARP poisoning or physical position, the data traveling over the network cannot be captured. A wireless LAN uses a physically public medium for data transfer; therefore, every packet traveling between a wireless client and an AP is transmitted via radio signals and can be captured by any client within radio range. Although captured data may be encrypted and not readily viewable, it should be noted that any client can obtain the ciphertext in some manner despite layer-2 or logical boundaries.

To address this issue the 802.11 standard provides a data confidentiality service. The legacy system provided the Wired Equivalent Privacy (WEP) protocol, which encrypted each message with a symmetric key before transmission. However, this protocol was successfully attacked and automated tools built to enable WEP keys to be cracked by anyone with basic IT skills and access to the freely available toolset.[7] Overtime, improvements have been made to the WEP protocol, moving it toward a more respectable enterprise solution, but these fell short and were soon superseded by newer, ground-up protocols. Additional controls are needed to protect the network if WEP must be used, such as with older existing equipment.

The RSN system provides two new data confidential protocols called *Temporal Key Integrity Protocol* (TKIP)[8] and *Counter Mode with Cipher-Block Chaining Message Authentication Code Protocol* (CCMP).[9] As well as confidentiality, both protocols provide message integrity as well.

33.2.4 Key Management. Secret-key encryption itself is a relatively simple process designed to protect data long enough for the encryption keys to be changed at a determined interval. Designers start by choosing a suitable proven algorithm, protocol, and key length that they are confident will protect user data for a defined

period. The user (or user program) then provides that algorithm with the plaintext data and an encryption key, and the data are then encrypted and ready for transmission. Due to the rapidly advancing speed of processing (of individual CPUs and CPUs running in parallel), no level of any encryption can provide definitive protection for an unlimited period of time—the time required for brute-force testing of all possible keys in the keyspace—necessitating the need for a key scheduling and management system.

Securely and repeatedly establishing mutual encryption keys is the single most complex issue plaguing cryptographic communications between two parties in physically separate locations. The repeatable process of creating keys, mutually (and securely) exchanging different keys or independently deriving the same key, and finally destroying them, is key management.

The legacy 802.11 algorithms did not provide any specific key-management functionality and vendors were left to design their own. The most common system in early standalone APs was to manually enter a defined-length static WEP key into each client and AP. For home users, the RSN system defines the manual entry of a common variable length PSK or passphrase in the client and AP. From this PSK, a key-management process derives working cryptographic keys, which are changed for each message.

For enterprises, RSN uses the 802.1/EAP framework to establish a secure channel during the user- and device-authentication phase, allowing a pairwise master key (PMK) to be set up between the client and the AP. From this PMK, a key-management system establishes working cryptographic keys which are changed for each message.

33.3 IEEE 802.11 ROBUST SECURITY NETWORK.

In June 2004, the IEEE released the 802.11i standard to improve the security of 802.11 networks. This new system is called the *Robust Security Network* (RSN) and is designed for both personal and enterprise users. Enterprise use is based on the 802.1X protocol to provide authentication and establish a security context. A "personal" profile uses a *pre-shared key* (PSK) based on a password provided for consumers and SOHO users who do not require the necessary 802.1X backend authentication infrastructure.

33.3.1 Features.

The core protocol of RSN is IEEE 802.1X, which forms *RSN Associations* (RSNA) with the wireless network. RSN provides the following features:

- Mutual authentication mechanisms. These mechanisms can authenticate users as well as the network client or machine. The AP and backend authentication server can also be authenticated to the client defeating rogue AP and man-in-the-middle attacks.
- Key-management algorithms
- Cryptographic key establishment (through PMK and Pairwise Transient Key [PTK] establishment)
- Cryptographic message integrity codes to defeat the bit-flipping attacks possible in the original standard (WEP)
- Two data privacy protocols which also implement message integrity:
  1. *Temporal Key Integrity Protocol* (TKIP), which is an optional protocol specifically designed so that existing WEP based hardware can be upgraded to use it.

2. *Counter Mode with CBC-MAC Protocol* (CCMP), which is mandatory for RSNA compliance. It uses the Advanced Encryption Standard (AES) in Counter mode for confidentiality and CBC-MAC for authentication and integrity. CCMP is a strong protocol that has been designed for the next generation of wireless equipment.

TKIP was designed as a temporary stopgap that would work on existing WEP-based hardware until new hardware containing the CCMP protocol became commonplace. Whereas TKIP continued to use the existing RC4 encryption algorithm that was at the heart of WEP, CCMP uses the AES algorithm and requires more powerful hardware—no longer a problem at present.

33.3.2 802.1X Overview. 802.1X was originally designed for port-based network-access controls for IEEE 802 LAN infrastructures. These infrastructures include Ethernet, token ring, and wireless networks. 802.1X authenticates and authorizes devices attached to an LAN port, and will not allow a device to access the network if authentication fails.

802.1X defines three roles:

1. **Authenticator**. The device that authenticates a network device before allowing it to access network resources. In an 802.11 BSS network the AP is the authenticator.
2. **Supplicant**. The device that wants to access network resources and needs to be authenticated.
3. **Authentication Server (AS)**. The AS performs the actual authentication of the supplicant on behalf of the authenticator. The AS can be located with the authenticator, but is commonly an external system such as a RADIUS server.

The 802.1X standard defines the object *Port Access Entity* (PAE) which operates the authentication algorithms and protocols in the supplicant and authenticator. An overview of the 802.1X architecture is shown in Exhibit 33.2 below.[10]

The authenticator has two logical ports; the first is an uncontrolled port that allows access to required functionality such as the authenticator PAE. The second port is the controlled port that allows access to the rest of the network. The status of the controlled port is set by the authenticator PAE and is dependent upon the outcome of the authentication between the supplicant and the authentication server.

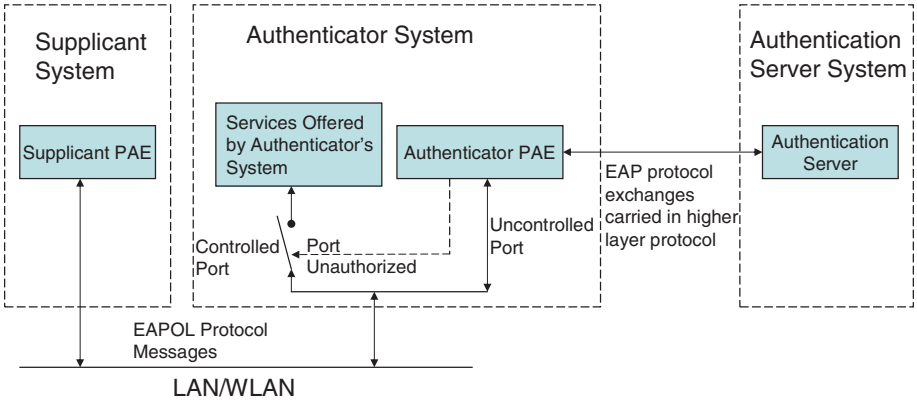


Exhibit 33.2 802.1X Architecture

The messages between the supplicant and authenticator use the Extensible Authentication Protocol (EAP) over LAN (EAPoL) framework defined in 802.1X. Communications between the authenticator and the AS leverage the EAP framework carried in a higher-layer protocol such as RADIUS.

33.3.3 EAP, EAPoL, and PEAP. 802.1X relies on EAP[11] to perform authentication of the supplicant and authenticator. The EAP protocol is a series of method interfaces that make up the framework known as EAP. EAP was originally designed for use on modem dial-up networks; therefore, the 802.1X specification details the expansion of EAP across Ethernet/Token Ring networks through the EAPoL (EAP Over LAN)[12] extension. In addition to placing EAP methods within an Ethernet payload, EAPoL specifies a number of additional functions to assist in the authentication process during discovery and key exchange. There are currently over 40+ different implementations of the EAP framework. The primary difference in EAP implementations centers on how the supplicant and authenticator are authenticated.

Standard EAP is not a mutual authentication protocol; only the supplicant is authenticated. This makes supplicants vulnerable to rogue AP attacks. Additionally, due to its original design for physical dial-up connections, EAP does not protect its authentication messages from eavesdropping. Therefore, the current EAP/EAPoL implementations establish secure tunnels to provide security prior to exchanging authenticating material. The following are the most common EAP/EAPoL enterprise 802.1X implementations ranked by the level of afforded security:

1. **EAP-TLS**[13]**:** This EAP implementation only allows mutual certificated-based authentication through *Transport Layer Security* (TLS) X509 certificates. Authentication of both the backend authenticating server and wireless client provides a strong level of wireless security, but requires a full enterprise *Public Key Infrastructure* (PKI) implementation to securely distribute and update client keys on a regular basis. An issue with EAP-TLS is most organizations do not have the necessary PKI to issue the supplicant client TLS certificates.
   - **EAP-TTLS** (*Tunneled TLS*)[14]**:** EAP-TTLS is similar to EAP-TLS and supports mutual certificate authentication but does not require client-side certificates. EAP-TTLS creates a TLS tunnel prior to starting any network authentication process and can therefore tunnel any password authentication mechanism, even insecure legacy mechanisms such as PAP.
2. **EAP-PEAP** (*Protected EAP*)[15]**:** In its native form, EAP-PEAP does not support mutual certificate-based authentication. Native EAP-PEAP uses TLS to authenticate the backend directory server only and leverages a password-based challenge-response process to authenticate the client. Later PEAP extensions help to mitigate this weakness. EAP-PEAP is native to most Microsoft Windows versions and is therefore very prevalent across the wireless industry. There are currently three primary types of EAP-PEAP, which provide varying levels of protection:
   - EAP-PEAP-MS-CHAP-v2 (*Microsoft Challenge Handshake Authentication Protocol*)[16]**:** The most common EAP-PEAP inner authentication method uses Microsoft's CHAP-v2 protocol to provide user-identifier (userID) and password challenge–response-based authentication. The MS-CHAP authentication process has historically allowed an attacker with physical access to the AP, to capture the provided challenge and challenge-response hash in plaintext. Due to recent cracking advancements, the time to crack an

MS-CHAP-v2 challenge-response hash has been reduced to days or less and should be avoided at all costs unless proper client supplicant configurations can be established.[17] This process is described in section 0 of this chapter.

- EAP-PEAP-TLS[18]**:** PEAP-TLS is the second PEAP inner protocol defined by Microsoft. PEAP-TLS tunnels the EAP-TLS protocol within PEAP to provide mutual X509 certificate-based authentication.
- EAP-PEAPv1 (EAP-GTC): The third implementation is defined by Cisco, which allows authentication using generic token cards such as RSA's SecurID token as well as user ID and password.

3. **EAP-LEAP**[19]**:** A proprietary protocol developed by Cisco, which performed challenge-response MS-CHAP-v2, based username/password authentication in cleartext. An attacker targeting a network employing EAP-LEAP only needs to monitor traffic to capture challenge-response hashes for offline dictionary attack. This is in contrast to EAP-PEAP-MS-CHAP-v2, which requires a rogue masquerading AP for an attacker to intercept MS-CHAP-v2 challenge-response hashes.
4. **EAP-FAST**[20]**:** (Flexible Authentication via Secure Tunneling) was developed by Cisco as a replacement to the vulnerable LEAP protocol. EAP-FAST introduced secure pre-authentication tunnels without using certificates. EAP-FAST has secure mutual authentication capabilities, but also has an automatic PAC provisioning option—which is vulnerable to man-in-the-middle attacks.

33.3.4 Insecure Legacy EAP protocols. When EAP was first integrated into 802.11, the EAP-MD5 and Cisco's EAP-LEAP protocols were commonly used. Unfortunately, these turned out to be vulnerable to attack, which spurred the development of the secure protocols covered above. EAP-MD5 and EAP-LEAP should not be used for enterprise wireless deployment.

33.3.5 Detailed EAP/EAPoL Procedure. Exhibit 33.3 shows the high-level RSN security association management flow. It consists of five stages that:

1. Establish a secure channel between the authenticator and authentication server
2. Locate the network, negotiate cryptographic algorithms, and associate to it
3. Provide 802.1X authentication to the authentication server
4. Provide Mutual Authentication and establish pairwise cryptographic keys
5. Establish Group/Multicast cryptographic keys

The following sections describe these stages. There are two aspects that vary the security association flow. The first is whether the wireless network contains an Access Point (a *Basic Service Set* or BSS) or whether it is an *independent* BSS (IBSS), also known as an *ad-hoc network*, which is a peer-to-peer network topology.[21] The second is whether the master cryptographic key is a global PSK or if it is established during the 802.1X authentication protocol. Any variations caused by BSS/IBSS and 802.1X/PSK are described within the sections.

Stage 1—Establish a Secure Channel between Authenticator and Authentication Server

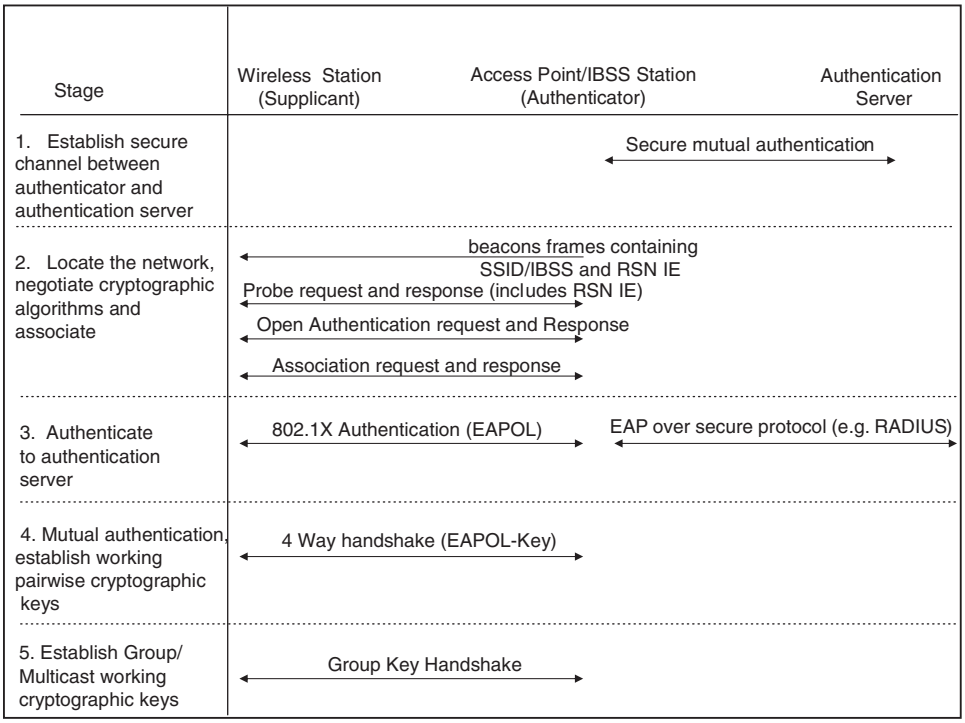


EXHIBIT 33.3 RSN Security Association Management

In this stage the authenticator and AS mutually authenticate one another and establish a secure channel between them using a protocol such as RADIUS, IP Security (IPSec), or TLS. This channel is used to securely carry the authentication exchanges between the supplicant and the AS. This stage is not required if the network uses a PSK.

Stage 2—Locating the Network and Associating To It

This stage is mostly the original 802.11 functionality for locating, authenticating, and associating to a wireless network. The key difference in RSN is the beacon frames, probe responses, and association requests contain Information Elements that indicate supported and available authentication and privacy protocols. In addition, the fast BSS transition protocol defined in 802.11r[22] enhanced these frames again to speed up AP roaming.

Stage 3—802.1X Authentication to the Authentication Server

The purpose of this stage is to mutually authenticate the supplicant and AS to one another and independently generate the PMK for use in stage 4. The EAP described above is used to achieve this. The messages exchanged between the supplicant and the AS are defined by the EAP method. An overview of this stage is given in Exhibit 33.4.

For a BSS network, the wireless client is the 802.1X supplicant and the AP is the authenticator. The AP relays the authentication messages between the supplicant and the authentication server, which can either be a separate service or built in.

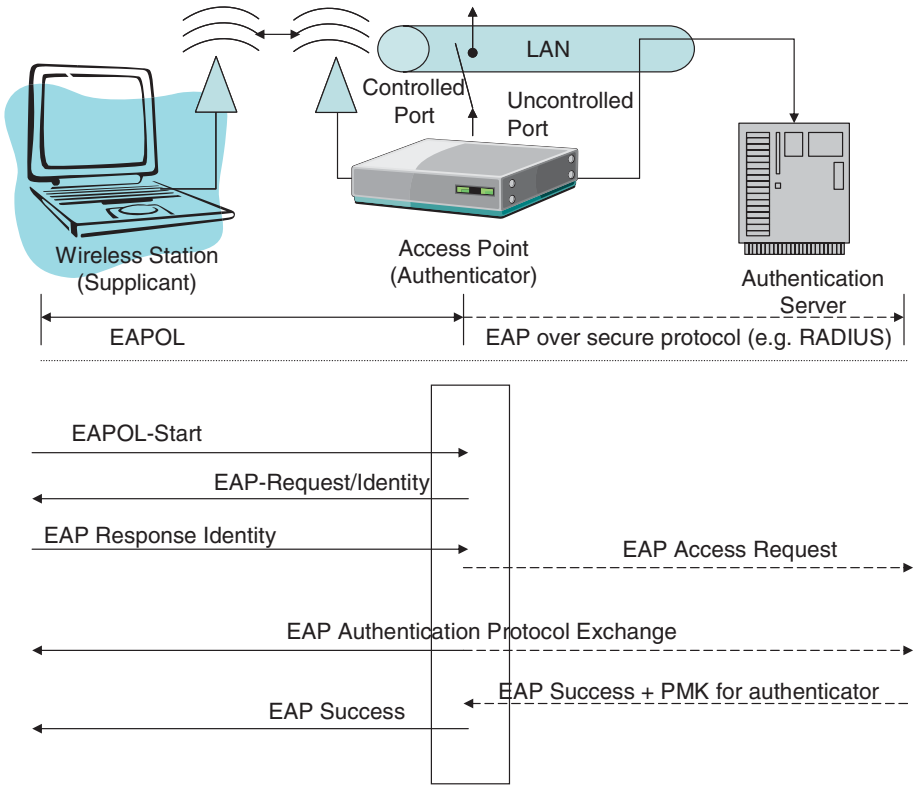


EXHIBIT 33.4 802.1X Authentication to Authentication Server

For an IBSS, the client wanting to associate to another client is the supplicant, and the target client is the authenticator. This means that clients in an IBSS can simultaneously be supplicants and authenticators depending on who initiated the association. Additionally, each IBSS client will need an authentication server unless the network uses PSKs.

As covered in stage 1, a secure channel has to be established between the authenticator and authentication server prior to this exchange; this serves two purposes. The first is to protect the integrity and authenticity of the authentication exchange, the second is to allow the AS to securely send the PMK to the authenticator once authentication is complete.

Stage 4—Mutual Authentication and Establish Pairwise Working Keys

Stage 3 established a PMK in both the supplicant and authenticator. If the network uses a *pre-shared key* (PSK) then the PSK is the PMK. This stage is called the 4-way handshake and has the following purposes:

1. To mutually authenticate the supplicant and authenticator to one another by confirming they both have the same valid PMK.
2. To generate a *Pairwise Transient Key* (PTK) from the PMK and fresh temporal keys bound to their MAC addresses.
3. To synchronize the installation of the keys in both devices.

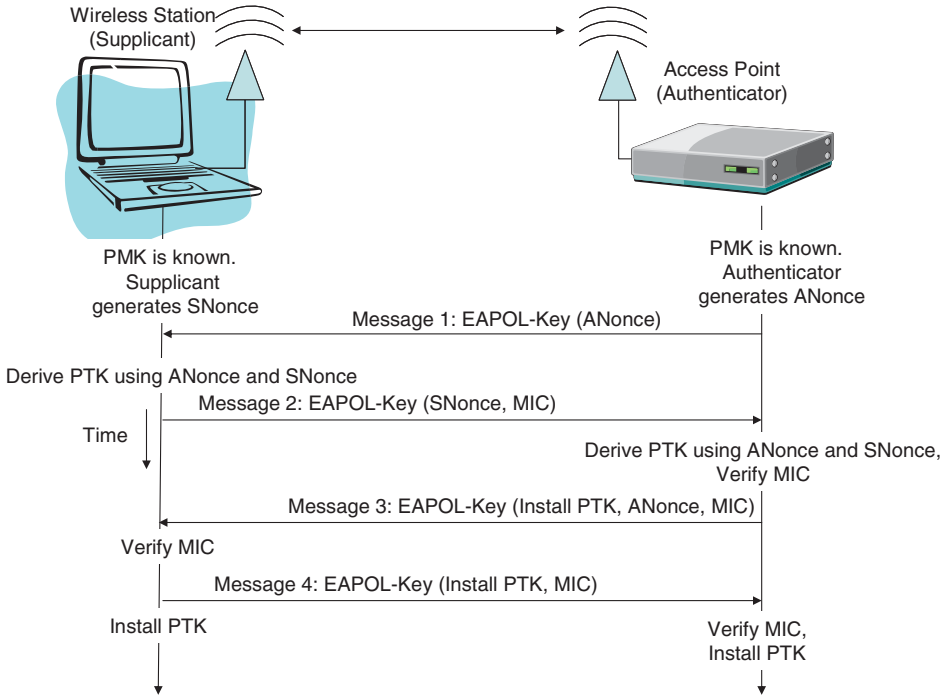


EXHIBIT 33.5 The 4-Way Handshake To Authenticate and Establish PTK

The 4-way handshake is implemented by EAPoL-Key messages exchanged between the supplicant and the authenticator and consists of the following as shown in Exhibit 33.5.

1. The authenticator and the supplicant both generate nonces for use in the authentication protocol. The authenticator's nonce is called ANonce and the supplicant's nonce is called SNonce.
2. The authenticator sends an EAPoL-Key message containing the ANonce (Message 1).
3. The supplicant derives the PTK using the ANonce and SNonce and calculates EAPoL-Key Encrypting Key (KEK) and EAPoL-Key Message Integrity Code (MIC).
4. The supplicant sends an EAPoL-Key message containing the SNonce and an MIC calculated using the EAPoL MIC Key (Message 2).
5. The authenticator can now derive the PTK because it has both the ANonce and SNonce. It then calculates the EAPoL-Key KEK and EAPoL-Key MIC Key and verifies the MIC in Message 2.
6. The authenticator sends a message containing the ANonce, and a flag instructing the supplicant to install the key. The message is also authenticated by an MIC (Message 3).
7. The supplicant verifies the MIC and sends a message to the authenticator confirming the installation of the key. This message is authenticated by an MIC and encrypted using the EAPoL-Key KEK (Message 4).

8. The authenticator installs the new keys and starts the last stage to establish the group keys.

Stage 5—Establish Group/Multicast Cryptographic Keys

In Stage 4, the PTK and temporal keys were established. These keys are used to secure an additional message pair in which the authenticator sends the group temporal key (GTK) in an encrypted form to the client. The group keys are used to secure broadcast messages such as ARP requests and multicast traffic.

In a BSS network, all clients have the same group/multicast key, which is sent to each client by the AP. However, for IBSS networks, there is no AP to set a common key, so each client has its own group transmit key that it sends to all clients in the IBSS. This is achieved by executing the 4-way handshake and group key handshake in both directions.

33.3.6 RSN Key Hierarchy and Management. A key-management system derives working (temporal) keys from a root master key. The exact key hierarchy varies slightly between TKIP and CCMP, but broadly follows the same system. This section describes the key-management system and notes the differences between TKIP and CCMP.

There are two key hierarchies: the first contains the pairwise keys, which are shared between two wireless devices (e.g., between two clients in an IBSS or between a client and an AP in a BSS). The second hierarchy is the group/multicast keys that are used for network broadcasts such as ARP requests or multicast traffic.

33.3.6.1 Pairwise Key Hierarchy. There is a maximum of four levels of cryptographic keys in the pairwise key hierarchy.

1. **The 802.1X Authentication Keys.** These keys only exist if the supplicant and authentication server mutually authenticate one another using preinstalled keys. An example is EAP-TLS, which requires both the supplicant and authentication server to have PKI credentials. The 802.1X keys are used to establish the pairwise master key (PMK).
2. **Pairwise Master Key.** The PMK is 256 bits and is either the key established in level 1 or a PSK installed in the devices. The PMK, device MAC addresses, and nonces are fed into a variable length *pseudorandom function* (PRF) to generate a pairwise transient key (PTK). The nonces are derived from a 256-bit key counter that is initialized at system start-up from a random number, the time, and the MAC address. The nonce is then incremented every time a key is changed.
3. **Pairwise Transient Key.** The PTK varies in length from 128 bits to 512 bits, and is split up into the required temporal keys. The key lengths and number of keys are dependent on the algorithm. For TKIP, the PTK is 512 bits and this is split into 5 temporal keys. For CCMP, the PTK is 384 bits and is split into 3 temporal keys.
4. **Temporal and per-packet keys.** The temporal keys are mixed with variable data such as packet counters, which has the effect of creating a fresh key for every packet. See Exhibits 33.6 and 33.7 for a description of the TKIP and CCMP temporal keys.

EXHIBIT 33.6 TKIP Temporal Keys

TKIP Temporal Keys		
Key Name	**Length**	**Purpose**
EAPoL-Key MIC Key	128 bits	This key is the first 128 bits of the PTK and authenticates the EAPoL-Key messages that establish the temporal keys. The key is used to calculate a message integrity code (MIC).
EAPoL-Key Encryption key	128 bits	This key is bits 129–256 of the PTK and is used to encrypt the contents of EAPoL-Key management messages.
Temporal Encryption Key	128 bits	This key is bits 257–384 of the PTK and is used to encrypt a packet using WEP.
Temporal MIC Key 1	64 bits	This key is bits 385–448 of the PTK and is used to authenticate messages in one direction.
Temporal MIC Key 2	64 bits	This key is bits 449–512 of the PTK and is used to authenticate messages in the other direction.

Exhibit 33.8 shows the entire key hierarchy for TKIP.

33.3.6.2 Group Key Hierarchy. The group key hierarchy is similarly structured to the pairwise key hierarchy. The authenticator creates a group master key. This master key, the authenticator's MAC address, and a *group nonce* (GNonce) are processed by the PRF to produce a group transient key (GTK), which is then split up into temporal keys.

33.3.7 Temporal Key Integrity Protocol (TKIP). The Temporal Key Integrity Protocol (TKIP) is a suite of algorithms to resolve the known key management issues in WEP for the existing 802.11 equipment already in the field. A key management mechanism was never specifically identified in the WEP protocol, which TKIP was designed to address.

To ensure compatibility with the existing base of wireless products, TKIP needed to take into account the hardware architecture of the existing wireless products.[23] Wireless products have two CPUs; the first is in the MAC chip that implements the wireless protocol, the second is the host CPU.

EXHIBIT 33.7 CCMP Temporal Keys

CCMP Temporal keys		
Key Name	**Length**	**Purpose**
EAPoL-Key MIC key	128 bits	This key is the first 128 bits of the PTK and is used to authenticate the EAPoL-Key management messages by calculating a MIC
EAPoL-Key Encryption key	128 bits	This key is bits 129–256 of the PTK and is used to encrypt the contents of EAPoL-Key management messages
Temporal Key 1	128 bits	This key is bits 257–384 of the PTK and is used for message encryption, authentication, and integrity. The nature of the CCM algorithm does not require a separate MIC key.

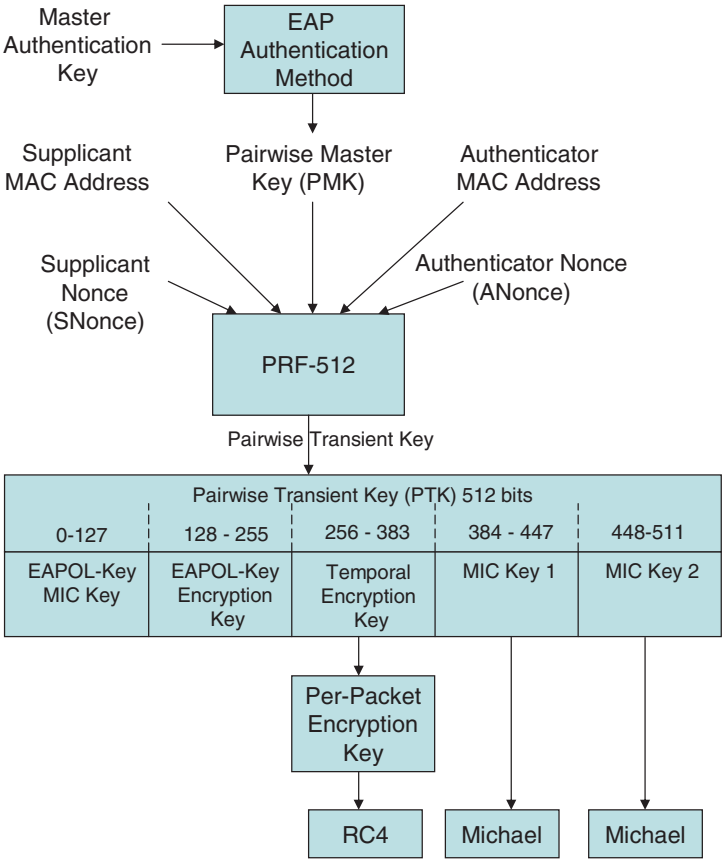


Exhibit 33.8 TKIP Key Hierarchy

For APs, the host CPU is a dedicated CPU that runs the AP. For wireless LAN PC cards the host CPU is that of the host PC. For performance reasons most MAC chips have a hardware RC4 encryption engine to perform the encryption. To avoid an unacceptable performance impact by moving encryption to the host CPU, TKIP needed to use this existing MAC encryption engine. This required TKIP to continue to use RC4 and WEP but in a way to overcome the identified problems.

The designers came up with a solution to the issues, and the result was TKIP with the following functionality:

- Each MAC service data unit (MSDU) is authenticated and integrity protected using a keyed cryptographic MIC. An MSDU is a message to be sent to another client. An MSDU may be fragmented into more than one MAC protocol data unit (MPDU), which are packets/frames sent by the physical layer. When all the MPDUs have been received, the recipient reconstructs the MSDU and passes it up the protocol stack. The source and destination addresses as well as the MSDU plaintext are included in the MIC calculation. This prevents forgery and masquerading attacks.
- However, the strength of the MIC is limited and can be compromised by trial and error. To overcome this shortcoming, TKIP provides optional countermeasures that stop communications for a period of time (60 seconds) when an invalid MIC

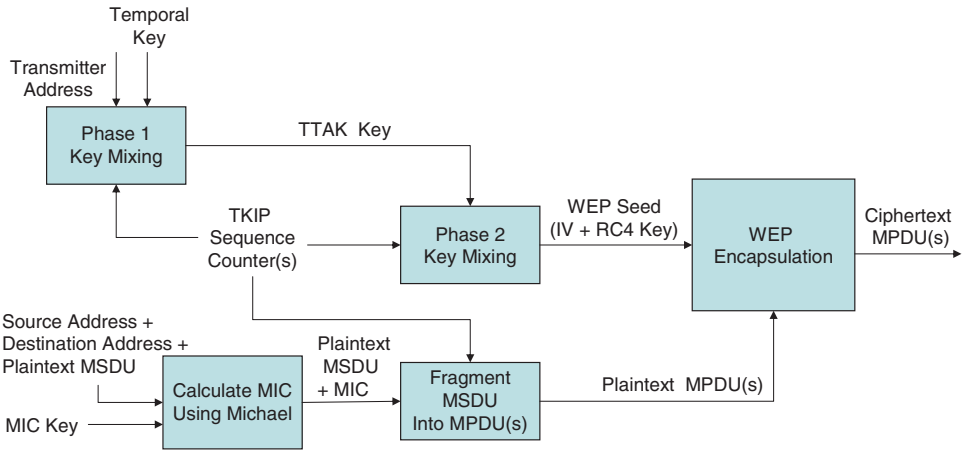


EXHIBIT 33.9 TKIP Algorithm in the Transmitting Client

is received, and then immediately forcing a re-key of all clients. The designers estimated that the MIC with countermeasures would resist attack for about one year before an attacker would correctly guess the MIC.

- Each TKIP MPDU (packet) has a sequence number encoded in the WEP Initialization Vector. Any MPDUs that arrive out of order are dropped.
- TKIP mixes a temporal key, the transmitter's address, and a sequence counter in a two-phase protocol to form the RC4 WEP seed. The mixing is done in a way to defeat weak-key attacks.
- RSNA uses 802.1X EAPoL-key message to regularly change the temporal keys so that RC4 key streams are not reused. The key change is triggered automatically when the sequence counter is close to exhaustion.

These controls largely address the issues identified in the legacy system by introducing cryptographic message authentication, preventing key stream reuse, and never using a weak IV. However, TKIP is still not considered a strong solution because of the strength of the MIC. TKIP should only be used as an interim measure until existing equipment can be replaced with CCMP-capable equipment. Exhibits 33.9 and 33.10 show the TKIP process in the transmitting and receiving clients.

33.3.8 Counter-Mode/CBC-MAC Protocol (CCMP). CCMP is the mandatory protocol defined in RSN to provide confidentiality, authentication, integrity, and replay protection for the next generation of wireless equipment. It is not possible to upgrade existing equipment to use this protocol due to resource requirements on hardware, which implements AES.

CCMP uses AES encryption in counter mode to provide confidentiality and CBC-MAC (AES-CCM) for message authentication and integrity. The inputs to the algorithm include:

- A 128-bit block cipher encryption key (the temporal key).
- A nonce based on an incrementing packet number that is used only once with the encryption key to encrypt a message. Reusing a nonce to encrypt more than one message destroys its security properties.

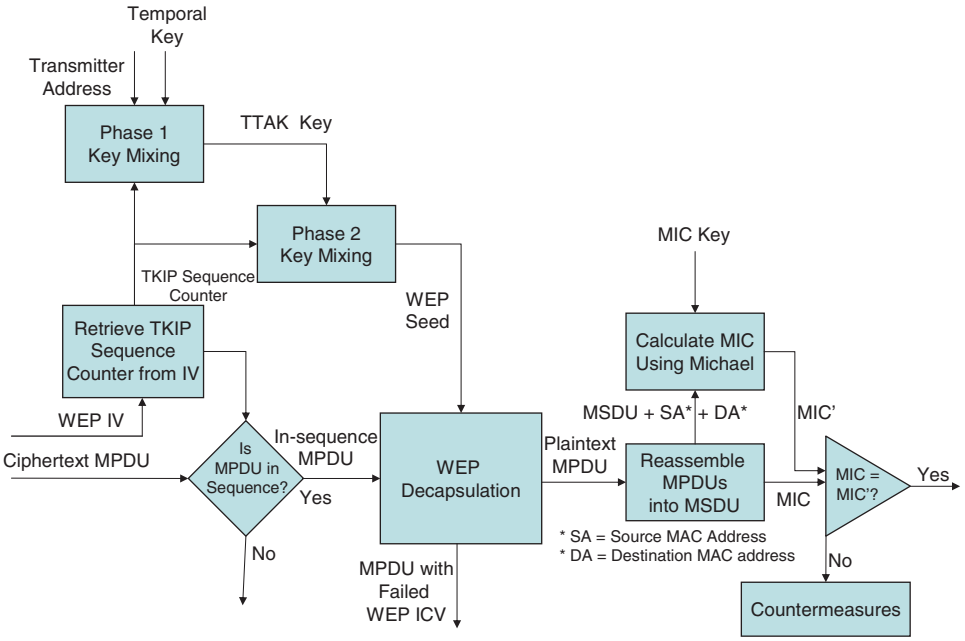


EXHIBIT 33.10 TKIP Algorithm in the Receiving Client

- The message to be encrypted.
- Additional data to be authenticated but not encrypted, such as the packet header containing the source and destination MAC addresses.

33.3.9 Fast Secure Roaming. After studying the RSN authentication system, the reader is likely to be struck as to the complexity of the system and the number of messages required to establish an association to the network. In an enterprise scenario using the 802.1X/EAP framework, roaming between APs could take up to one second to complete. This creates usability issues for time-sensitive applications such as mobile voice.

RSN contains the PMK caching and preauthentication protocols to try and address this problem; however, both protocols have limitations. PMK caching involves the client and APs remembering the keys they used with each other when the client roams away. But this is only useful when roaming back to APs you have roamed from. Preauthentication addresses this issue by completing the 802.1X authentication and establishing keys with APs as they come across them, even if the client never roams to that AP. This creates significant unnecessary authentication traffic that could overwhelm the authentication server.

The 802.11 committee decided to look at the problem again and developed the 802.11r standard that enhanced RSN PMK caching to establish a cached PMK with a central controller, which then distributes a variant of the PMK to APs as needed. This protocol is called Opportunistic Key Caching (OKC) and means that the lengthy 802.1X authentication is no longer required, reducing roaming times to approximately 100 milliseconds. 802.11r also enhanced the authentication and association messages to include the OKC parameters saving the need for additional message pairs.

33.4 FUNDAMENTAL WIRELESS THREATS. It's important to ensure an enterprise is prepared to manage the many fundamental threats which are enabled by the existence of a wireless medium. As an enterprise examines their infrastructure and needs as a whole, there is an inherent lack of consideration for the increased threat exposure due to the elimination of physical layer controls. These types of controls are typically taken for granted and not considered despite the change to a completely new physical medium. It is important to understand the fundamental differences between a wired and wireless medium and be prepared to take action through implementing necessary controls to secure a wireless network to a level equivalent to a wired LAN.

The implementation of a wireless network expands network availability beyond that of facility walls or even parking lots. It's not uncommon for adversaries to sit within a corporate lot and use special antennas (called *cantennas*) to eavesdrop on wireless communications from a safe distance. Wireless networks rely primarily on logical controls which occur at data link layer and above in the form of authentication and encryption to compensate for the fundamentally public availability of the medium.

Driven by the ease of accessibility of wireless communications, the threat model an organization faces due to the existence and use of a wireless network may change as previous threats are eliminated (such as unauthorized network jack use) and others are created. The increasingly complex threat landscape and security-monitoring responsibilities can affect an enterprise's decision to implement a wireless solution for a particular use or limit its scope to designated functional uses.

Within this section, we will cover the fundamental threats enabled by the use of a wireless network and we will take a deeper dive into a subset of specific technical threats, which should be a priority consideration for an enterprise. It's not necessary for an enterprise to understand every threat they may face, but it is necessary to acknowledge their presence and take actionable steps toward countering these threats.

33.4.1 Unauthorized Network Extensions. When transitioning to a wireless environment from a wired network or a combination of the two, threats which were seemingly easy to identify in the past are now increasingly difficult to manage. Due to the inherent lack of transparency in regulated and unregulated airwaves, it is possible for a nonauthorized access point (AP) to hide amongst authorized APs through SSID spoofing. In a wired environment, it is more difficult to hide an unauthorized device due to physical appearance. An enterprise's ability to identify an unauthorized AP will require additional controls to perform this operation after authorized wireless access has been deployed.

33.4.1.1 Rogue Access Points. A rogue AP is attached to the network without authorization or which impersonates an attached device. Rogue APs come in multiple forms:

- **Nonmalicious Internal Rogue AP:** This type of rogue AP is commonly implemented by an employee or staff without malicious intent, but may be abused by an outside attacker. By connecting a Small Office Home Office (SOHO) grade AP to the internal network, enterprise grade protection mechanisms designed to keep attackers out, such as perimeter firewalls, can be bypassed with direct access to the internal LAN. Endpoint security mechanisms are also invalidated, such as host-based antivirus and firewalls, due to the Network Address Translation (NAT) capability of an SOHO router. Consumer-grade security controls present

on an SOHO AP, combined with the lack of proper security configuration by the installer, can lend themselves to successful compromise, which would result in direct access to the wired internal network. Nonmalicious rogue APs can typically be identified through a nonenterprise standard SSID such as "Linksys" or "netgear."

- **Attacker Malicious Internal Rogue AP:** An attacker can masquerade an AP as a legitimate enterprise AP by using the same SSID as authorized enterprise APs. In this situation an unsuspecting endpoint client will automatically connect to any AP, which broadcasts the requested SSID (such as "employeeWireless") and will select the AP with the strongest signal. An attacker can then perform denial of service attacks to force clients onto their malicious AP.

  A malicious rogue AP with internal access can be leveraged by an attacker with physical access to the facility to capture credentials or other data from unsuspecting endpoint clients, which leverage the AP to access internal resources such as administrative panels or employee portals. An attacker with an internally connected AP has multiple benefits as they are able to seamlessly allow a user to access all regularly accessible resources. This type of attack would require physical access to the facility and corresponding network jacks.
- **Attacker Malicious External Rogue AP:** An external rogue AP uses similar tactics to an internal rogue AP to force unsuspecting endpoint clients to connect to it. The primary difference is the physical proximity of the AP and the network it provides access to. It is much easier for an attacker to deploy an external rogue AP in a nearby corporate parking lot than gain access to the facility and plant an internally connected AP. The downside of an external rogue AP is the inability to provide requested internal resources to a connected client. If the desired resource is the Internet, this can be provided through a forwarding process allowing the attacker to view all traffic.

  Both types of malicious APs allow an attacker to capture sensitive network traffic and attack the endpoint itself by attempting to exploit vulnerabilities in the endpoint's operating system or network services. Client configurations typically mandate that a client wireless radio connect to the strongest signal strength that matches the desires SSID and Authentication requirements; e.g., WPA/WPA2-PSK, WPA/WPA2-Enterprise (EAP-PEAP, EAP-TLS).

Threats enabled by rogue APs include:

- An attacker's gaining access to the wired network by compromising weak or absent security controls on a nonmalicious employee-installed rogue AP
- An attacker's obtaining sensitive data through network captures of clients connected to an attacker-provided rogue AP (such as domain credentials). This threat is usually more effective if used in conjunction with social engineering tactics, such as injecting IFRAMES within HTTP traffic that directs them to a malicious site, which contains browser exploits for various plugins such as ActiveX or Adobe.
- An attacker's gaining access and penetration testing enterprise endpoint machines connected to an attacker-provided rogue AP at a different physical location.
- Internal users' bypassing wired 802.1X port security configurations through the use of NAT, which allows multiple users to communicate through the WAN port of an installed SOHO AP.

33.4.2 Layer 1 Attacks. Wireless mediums can be heavily relied upon by staff or business applications to provide connectivity to wired network resources in locations where physical connectivity is difficult to implement. The absence of an enclosed private physical medium allows for tampering of the layer 1 wireless carrier signal regardless of upper layer protective controls. This is a fundamental threat to any form of wireless communication and not necessarily specific to 802.11, although each wireless protocol such as CDMA, 3G, or 802.11 may include a mechanism to detect tampering, they cannot prevent it from occurring.

- The 802.11 protocol functions within 2.4 GHz band. The 2.4 Ghz consumer band can be interfered with by using something as simple as a microwave oven
  - The 2.4 GHz spectrum only provides three nonoverlapping channels. The 802.11 spectrum runs between the range of 2.412 GHz and 2.462 GHz. Overlapping channels are side-by-side channels, which through propagation and other wireless tendencies, can bleed into the neighboring 20 MHz bandwidth allocation. A nearby AP can cause a minor denial of service condition by using an overlapping channel.
- Management frame spoofing
  - 802.11 management frames are part of the MAC layer of the 802.11 protocol. Management frames are used to bootstrap the authentication and association of clients to an AP. Management frames "manage" the connection between an AP and all connected clients. There are various types of management frames, some of which are received by any client within range of the AP, and others which are targeted at connected/connecting clients.

Most significant management frame types:

- Probe Request/Response
- Authentication
- Association Request
- Association Response
- Dissociation

As with a wired LAN, the unauthenticated nature of initial data link layer protocols, such as 802.11 management frames, can allow an attacker to spoof and interfere with an existing session or the session initiation process. This fundamental issue is comparable to the unauthenticated Ethernet → IP ARP protocol, still vulnerable over 802.11, but easily abused across a wired connection.

The security afforded to wired bootstrap protocols relies on their physical controls, none of which are afforded to wireless. This void prompted IEEE 802.11 family of standards to adopt the 802.11w specification. This specification details the use of MICs to authenticate management frames to their source. This would effectively eliminate the ability for an unauthenticated attacker to force clients off of a legitimate AP to an attacker provided rogue AP. 802.11w sounds great on paper, but in reality, implementation will take years. The process of authenticating material for each management frame increases the load on an AP's already busy processors, which means a software upgrade may not be feasible for all devices. In addition, all client devices must support 802.11w to create outgoing and process incoming MIC data. Client support would require new

driver firmware for wireless radios and potentially more hardware acceleration for potentially all connected clients depending on the AP's implementation of 802.11w.

33.4.3 Endpoint Attacks. A new class of endpoint attacks is enabled by wireless endpoints. Endpoints can include anything from a laptop, wireless printer, or handheld scanners. In a pure wired environment the threat of wireless endpoint compromise is slightly reduced due to the absence of corporate wireless authenticating materials stored within the endpoint's wireless supplicant. Authenticating materials include items such as domain username/password pairs, certificates, or even pre-shared keys. Storing authenticating materials on an endpoint means the "keys to the kingdom" can be obtained and used to access the network without even entering the facility.

While the ability for an attacker to compromise an endpoint with a hotspot style attack, is still a viable threat to the integrity of an endpoint, endpoint security measures such as host-based firewalls, IDS, and regular patching can mitigate this threat to an acceptable level. It is necessary to understand that an endpoint is still susceptible to the same attacks in an enterprise wireless environment, but the outcome of these attacks could result in the capture of authenticating materials used to access the corporate wireless. The capture of authenticating materials avoids the need to compromise the endpoint to gain internal access to a corporate network. If employees do not use wireless enterprise access, they are less likely to succumb to social engineering attacks combined with a hotspot attack. An example of this situation involves an attacker using a hotspot attack to gain connectivity to the endpoint, intercepting their Internet browsing, and prompting the user with a spoofed corporate intranet site, which requests remote VPN credentials through a simple Web form.

Endpoint attacks enabled by enterprise wireless access:

- Employer identification through descriptive SSID probe requests.
- Increased effectiveness of Hotspot + Social Engineering attacks based on the fact that a user regularly accesses the enterprise wireless, in contrast to a user whom has never had wireless access.
- RADIUS impersonation for capturing WPA/WPA2-enterprise EAP credentials.

33.5 SPECIFIC WIRELESS SECURITY ATTACKS

33.5.1 MAC Address and IP Spoofing. MAC address–based access control mechanisms are the lowest hanging fruit of wireless security threats. This form of access control is not enterprise grade and primarily originated in SOHO grade APs and hotspots for tracking authorized connections without issuing unique keying material. Pure MAC address authentication is based solely on the connecting devices MAC address; if the identified MAC address is included in a predefined list of allowed MAC addresses, the device is allowed to associate. This is obviously not a scalable means of authentication, due to constant updates in the event of static mappings, and is commonly used to track valid client connections after initial Web-based authentication.

This type of situation is commonly referred to as a Captive Portal or hotspot-based authentication. The AP operates in open authentication mode, allowing anyone to connect to it, but users who have yet to authenticate cannot access the Internet and are dropped into a restricted penalty box. Users are prompted with a Web interface to enter credentials and, upon successful network-based authentication, gain access to the Internet or other protected resources. Captive Portal authentication takes place at

layer 4 (Application) and relies upon layer 3 (IP address) and layer 2 (MAC address) as reliable mechanisms to track users which have successfully authenticated.

This attack is enabled by the fact that all traffic in a Captive Portal system is transferred unencrypted. No encryption necessary to allow clients to access the Web-based authentication page with no previously established encryption keys or previous knowledge of the AP. An attacker now has the ability to capture plaintext traffic occurring on the AP to and from each client. This is performed through setting a compatible wireless card into monitor mode. Monitor mode allows an attacker to see all traffic on a specific 802.11 channel without needing to associate with an AP. On wireless networks which deploy an encryption mechanism such as WPA/WPA2-Enterprise, an attacker will see no plaintext payloads for 802.11 frames.

An attacker can leverage this to monitor traffic and identify MAC address and IP addresses of connected/authenticated clients. This information can then be spoofed by an attacker to gain access to the portal's restricted content (Internet or other resources) without being directed by the portal to the authentication page. Authenticated clients are maintained by the AP via MAC address, IP address, or a combination of both.

33.5.2 EAP (WEP/WPA/WPA-Enterprise) RADIUS Impersonation Attack.

Wireless implementations, which leverage RSN with clients which have not been configured to trust a specific AP type, and do not employ client mutual certificate-based authentication, are inherently vulnerable to a deadly combination of rogue AP and authentication server (e.g., RADIUS) spoofing. This attack relies on the ability for an attacker to force clients to connect to a malicious AP and begin the expected EAP authentication method. In a typical EAP RSN deployment, in response to an authentication request a RADIUS server will send back a RADIUS access challenge, which is forwarded by the AP to the client. In this attack, the malicious AP is configured to use an attacker provided RADIUS server. This RADIUS server carries out the EAP authentication process and forces the client into providing their password hash in response to a fixed, known, challenge. Depending on the EAP implementation in use, the username associated with the client may be provided in plaintext within the EAP-Identity-Response sent prior to authenticating material exchange.

Since the attacking rogue AP uses a fixed challenge we now possess the ability to brute force the password using standard dictionary attacks. This attack only applies to EAP implementations, which do not use client certificates for authentication, but leverage username and password combinations, such as EAP-PEAP with MS-CHAP-v2.

If a client connects to a rogue AP, the spoofed authentication server certificate will likely not validate correctly and the user will either not see anything or be prompted with a small message. See Section 6.2 for additional information pertaining to proper client supplicant configuration for server-side certificate verification.

EAP challenge/response hashes can be captured and cracked using the following tools:

- **FreeRADIUS WPE (Wireless Pwn Edition):** Used to readily deploy a RADIUS server to accept connections from the rogue AP using a spoofed version of the expected EAP type.
- **AsLEAP:** Used to crack EAP (PEAP, LEAP, etc.) challenge/response hashes using the password hash, known challenge, and word dictionary as input.

33.5.3 Management Frames. 802.11 frames carry protocol and data for upper layer protocols. Primarily, two types of 802.11 frames are commonly used, data and management frames. Data frames are used as carriers for upper layer protocols, while management frames carry specific information regarding link operations. Historically, management frames are a fundamental weakness to 802.11 security. Across the majority of wireless implementations management frames are never secured and therefore always accessible despite authentication or encryption mechanisms, which may secure the payload of a data frame's contents. As described previously in section 4.1, 802.11w attempts to address this issue through the use of MICs, but like most wireless improvements. the industry lags behind fully adopting these controls due to costly client upgrades and new AP devices and/or firmware updates. The lack of adoption is likely due to the fact that management frame spoofing is not readily seen as a direct threat, despite its deadly impact when used in combination with other attacks.

Management frame spoofing and manipulation enables the following types of threats:

- **Beacon frame spoofing:** Allows an attacker to masquerade as any AP. Crafting a beacon frame will make you show up in a network client "scan" of available networks. This vulnerability enables hotspot and RADIUS impersonation attacks to take place and is largely unstoppable due to the necessity for unauthenticated beacon frames to be accepted by any unassociated client within range.
- **Authentication/Deauthentication:** These frames are heavily used by 802.11 during the secure authentication process. Used to relay the intention of a client to connect to an AP. Even in situations where an AP is using "open" authentication and requires no keying material an authentication frame procedure is performed. In the event of WEP pre-shared key authentication these frames are used to carry traffic associated with this network-based authentication.
- **Association/Reassociation/Disassociation**: Association response frames are used to synchronize radios between an AP and a client radio. This includes details such as supported data rates. This process allows the AP to allocate memory for the client and begin processing communications. Reassociation frames are used when a client roams to another AP within the same BSSID, this indicates to the original AP that it should forward any remaining buffered frames. A dissociation frame is sent by the client to gracefully remove itself from an AP and allow the AP to garbage collect any necessary memory.
- **Request to Send (RST) and Clear To Send (CST**): RST frames are optional and provide a way to survive in a condition where one client cannot see another client's traffic due to physical proximity to the AP. This frame is sent to the AP and the AP will response with a CTS frame. The CTS frame includes a narrow time frame for the client to send data where all other connected clients must wait. This procedure reduces collisions from hidden AP clients.[24]

The manipulation of management frames occurs by injecting them into the airwaves and spoofing the source as if they came from the AP BSSID. These frames are not viewable from a standard packet capture on the host system, as the wireless card driver will not pass this traffic to the underlying operating system. Spoofing management frames does not affect an AP, as it is unaware of their presence.

The following Aircrack[25] tools can be used with a compatible wireless card set into monitor mode to manipulate management frames[26]:

- **Aireplay-ng -0:** Spoof deauthentication to remove a client from the AP by sending dissociate packets to one or more clients, which are currently associated by forging a disassociate packet with the source MAC address of the target.
- **Aireplay-ng -1:** Spoof authentication with an AP using a specified source and destination MAC address by sending association management frames.

33.5.4 Hotspot and Cafe Latte Attacks. Depending on the wireless chipset and driver, wireless clients typically actively probe for all networks they have associated with in the past and when an AP responds they will automatically associate with it. The root of this vulnerability stems from the behavior of wireless clients and their lack of AP authentication during the early stages of authentication. Only the SSID and authentication/encryption type must match the "known" state for the client to begin the association process. Due to thin AP wireless deployments, a wireless client cannot be specific or remember the MAC address of the AP it requests to connect to, even though this would greatly curb this attack.

A wireless client probing for an SSID last used from a public Wi-Fi hotspot can be identified by an attacker and masqueraded as if they were the public Wi-Fi hotspot, even though the victim is nowhere within the location of the previously remembered public hotspot. Since most public hotspots don't require authentication, it's safe for an attacker to assume they only need to match the requested SSID and no other traits to facilitate a successful connection by the targeted victim. If the AP a client is probing for requires WEP or WPA, it will not automatically connect to any AP with the same name. The data transfer rates and encryption levels must align correctly, as described within the process of association.

An attacker sniffing the wireless for SSID probe requests will not necessarily know which networks require encryption and which don't, but trial and error will eventually yield the correct answer.

The following is the process associated with hotspot spoofing where an attacker recognizes a particular wireless client would like to connect to an AP, that may or may not actually be present (this is the nature of Wi-Fi autoconnect), and actively impersonates the requested AP to satisfy the client and gain momentary connectivity (in the event the client disconnects due to an invalid key management process).

33.5.4.1 Hotspot Attack Procedure

1. Victim machine sends SSID probe requests for APs it has communicated with in the past (supplicants keep this history to enable fast connections).
2. Attacking machine monitors airwaves for SSID probes and takes note of requested AP names (e.g., "coffeeshop," "airportWifi," etc.).
3. Attacking machine creates an AP with the SSID name requested by the client (e.g., "coffeshop," "airportWifi," etc.).
4. Victim machine automatically associates with new AP and encryption/authentication levels are accurate.

   At this point it is up to the attacker to correctly guess the encryption the client expects (WPA/WPA2/WEP) or open (none). It is very common for

Microsoft Windows and third-party supplicants to autoconnect to a known AP in the background and show no warning of the connection with the exception of a "connected" symbol in the taskbar.

5. The victim now has a functional connection to the attacker's AP and the attacker can establish *man-in-the-middle* (MITM) traffic by relaying it to the legitimate AP on a neighboring interface or the Internet, or attack the victim by exploiting vulnerabilities on the connected system.

An older yet valuable extension to this attack revolves around the SSID, which is broadcasted by a "parked" wireless card and is prevalent in older Windows wireless drivers. During Windows Wireless Zero Configuration, a wireless card will attempt to connect to all previously known networks before placing itself in a "parked" state. The "parked" condition consists of the adapter ceasing to attempt connections and using a dynamically generated SSID as a default value. Once this default value is assigned, the wireless card will no longer actively seek networks but will continue its normal beacon behavior, this time with the parked SSID. This behavior introduced a new vulnerability, which allowed an attacker to always have an SSID to rely on, even if the client has no previous networks to seek for. This vulnerability has since been addressed through setting an encryption level (WPA/WPA2) and random key for use by the wireless card during park mode.[27]

33.5.4.2 The WEP Cafe-Latte Factor. An added twist to hotspot spoofing takes it a step further by actively attacking a client that broadcasts a particular SSID from a network which deploys WEP encryption. This attack leverages a multitude of WEP vulnerabilities to crack a WEP key passed by a client. All clients will store the WEP key within the supplicant or registry of some sort for later use. An attacker can force the client to encrypt certain packets using this key while performing the aforementioned hotspot spoofing attacks to draw the client onto an AP that is masquerading as a requested one employing WEP encryption (albeit with an incorrect matching key). Specific packets are always sent by a client upon associating and during initial authentication as part of "quick reconnect," which are encrypted using the client's stored WEP key. These include ARP, DHCP, and specific 802.11 management frames.

Of particular interest are gratuitous ARP packets sent by a client upon associating to an AP it believes is legitimate. Through bit flipping techniques an attacker can forge encrypted ARP requests to a connecting client by reusing previously captured gratuitous ARP packets and turning them into a request. This technique leverages WEP's absence of integrity checking and reliance on a CRC error-checking mechanism as the sole integrity control. CRC algorithms are not intended to provide anything outside of error-based integrity and prone to abuse across WEP attacks. The compilation of hotspotting and WEP vulnerabilities can yield the WEP key to a corporate network within six minutes, without even needing access to the actual network. An attacker can now take his latte, your corporate WEP key, locate the corporate network, and connect from a parking lot.

33.5.4.3 Caffe-Latte Attack Procedure[28]

1. Victim machine sends SSID Probe requests for APs it has communicated with in the past (almost all supplicants keep this history to enable fast reconnections).

2. Attacking machine monitors airwaves for SSID probes and takes note of requested AP names (e.g., "coffeeshop," "airportWifi," etc.).
3. Attacking machine creates an AP with the SSID name requested by the client (e.g., "coffeeshop," "airportWifi," etc.).
4. Victim's machine attempts to connect but fails, likely due to an encryption mismatch.
5. Attacking machine modifies AP to possess WEP encryption with a false key.
6. Victim machine attempts to connect to the AP again and successfully authenticates due to a properly matching SSID name and encryption grade.
7. Victim machine pauses during further network-based authentication due to a key mismatch and never fully associates.
8. Attacker captures gratuitous ARP, DHCP, and other packets sent by the victim after association, which are encrypted with the victim's stored WEP key.
9. Attacker provides victim-encrypted gratuitous ARP packet to./cafe-latte resulting in bit-flipping to produce a CRC-correct encrypted ARP request.
10. Attacker replays encrypted ARP request to solicit responses from victim with new IVs until enough IVs are captured to exhaust 40-bit key space.
11. Attacker executes Aireplay-ng (or comparable WEP cracking utility) providing IVs captured during ARP request replays.
12. Attack receives plaintext hex of WEP key and uses to gain access to target network.

The following tools can be used to perform a hotspot and cafe latte–style attack:

- **Airbase-ng:** Create an AP with specified SSID
- **Cafe-latte:** WEP cracking script, which takes the captured gratuitous ARP packet and bit-flips it into a request that can be repeated to the client thousands of time.

33.5.4.4 What about WPA/WPA2? WPA and WPA2 implement cryptographic integrity checks, which remove any ability to deduce a plaintext key, but still suffer from a hotspot style of attack. A WPA/WPA2 tailored attack involves masquerading as a known AP where a client expects WPA or WPA2 encryption. Upon providing the correct encryption level the client will try to connect and fail due to the AP not possessing the correct key. Failure is expected, as the pre-shared key hashes with a salt of the SSID has already been sent from the client and captured by the attacker. This hash can then be brute forced.

33.5.5 WEP. Despite the public and lengthy exposure of WEP vulnerabilities and little security it actually affords, its widespread use at the enterprise level is shocking yet usually unknown due to its use in many small crevices of the business that it may have wedged itself in years ago.

WEP was conceived during a time when computing power, and in particular embedded computing power, was hard to come by. The computing power required to carry out the WEP RC4 encryption and key management is very minimal, which coincided perfectly with the computing power of many handheld and other embedded devices. This very principal is still the reason why WEP remains in many environments and will for the foreseeable future. Sunk costs in infrastructure devices incapable of performing

resource heavy functions such as AES and CCMP in WPA2. WPA's implementation of TKIP was precisely done to enable backward compatibility with WEP hardware. In today's enterprise environment it is commonplace to see WEP fully expunged from employee WLANs but prevalent in retail or distribution centers on handhelds and other portable devices.

Further compounding WEP's poor implementation of the RC4 stream cipher is the ambiguous definition within the standard for specific encryption function processes, largely leaving the decision to the vendor. This is particularly obvious in the way Initialization Vectors are generated; some vendors use a Pseudo-Random-Number-Generator (PRNG) and others merely start at 000000 and work their way to FFFFFF. We will briefly dive into the following tried-and-true core vulnerabilities affecting WEP, and its failure to provide proper key management, integrity, or encryption, and finish with a brief description of WPA/WPA2's advancements in this area.

33.5.5.1 (WEP) Repeating Initialization Vectors (IVs). WEP is home to a number of vulnerabilities including FMS, KoreK, PTW, and ChopChop. We are only going to cover the most famous and original attack, called FMS. WEP has been sunset ages ago in favor of alternative technologies; as such we will dive into the details of this vulnerability only to shed light on the advancements that have been made since its discovery.

WEP's underlying encryption revolves around the RC4 stream cipher. An initialization vector (IV) is used to ensure two identical inputs do not result in the same cipher text output. The IV is leveraged to lengthen the key life by uniquely generating an IV for each frame. WEP keys come in two forms depending on the version, an effective 40 bits (64 including IVs) with some implementations going as high as an effective 104 bits (128 including IVs) WEP uses a 24bit IV which, on a busy network, there is a 50 percent probability that the IV will repeat after 5,000 packets.

A fundamental solution to the WEP IV key scheduling vulnerability is the use of session keys established at the onset of authentication and encryption, which are used for a finite period of time during the encryption session. The process used to derive the temporary or session key is crucial to the success of a key scheduling process. WEP suffered from concatenation of RC4 output with an IV, and was improved upon with TKIP's "mixing" of these values prior to RC4. If a solid key scheduling and management function can be established, the encryption routine itself will deliver as expected. Attacks which target vulnerabilities in WEP's key scheduling process include the famous FMS attack, named after its inventors Fluhrer, Mantin, and Shamir.[29]

33.5.5.2 (WEP) Key Management. WEP does not provide any form of key management, such as the ability to rotate keys with clients. The result of no key management structure is the reuse of stale keys as input into the key scheduling and underlying encryption process for an extended period of time, further compounding the repeating IV vulnerability mentioned earlier. WPA and WPA2 provide session-based key management through the implementation of TKIP (WPA) and later CCMP/AES (WPA2). It is important to outline that the secret key in a pre-shared key implementation (WPA/WPA2-PSK) is only really leveraged to provided authentication; encryption is derived from the client and AP possessing the same key, but the pre-shared key is not directly involved. This is in contrast to WEP's use of a pre-shared key directly in the encryption process by leveraging it for authentication and including it in the RC4 stream cipher. Through the use of EAP and RADIUS, WPA and WPA2 Enterprise editions fully support key management through directory-based structures. These directories manage authentication keys outside of the WPA/WPA2 implementation and add increased

security through the combination of directory-based authentication key management (LDAP/Active Directory) and TKIP/CCMP encryption key management.

Tools used to exploit WEP vulnerabilities:

- **Airmon-ng**[30]: Used to place the wireless adapter into monitor mode.
- **Airodump-ng:** Used to capture WEP IVs using monitor mode.
- **Aireplay-ng-1:** Fake authentication to the AP; this will allow you to associate to the AP but not continue the authentication process. This is necessary for the AP to receive data from your wireless adapter.
- **Aireplay-ng-3:** ARP request replay attack mode, this will listen for encrypted gratuitous ARP packets sent by a legitimate connected client. This mode identifies ARP packets by size and other factors despite being encrypted. Once an ARP packet is found, Aireplay will bitflip and turn it into a request which is replied thousands of times. The purpose of this is to generate lots of encrypted packets from the victim and start exhausting IV space.
- **Aircrack-ng-b:** Performs statistical analysis and brute force of captured WEP traffic to recover the key. The process is to use statistical analysis to zone in on potential byte values at each location of the key, then use brute force to complete the process.

33.5.6 WPA/WPA2 Pre-Shared Key. WPA and later WPA2 aimed to address many of the deficiencies in WEP and WEP advancements. The following are attack vectors and mitigating controls implemented in the WPA standard:

- **48-bit Initialization Vectors:** Addresses WEP's 24-bit vulnerability and poor RC4 implementation through the use of TKIP/CCMP to significantly reduce IV reuse and improve the ability to "seed" the encryption function.
- **Key Management:** TKIP and CCMP/AES provide per-frame unique encryption keys in contrast to stale WEP key storage within client supplicants and reliance on the IV as a primary differentiating factor. WPA/WPA2 also natively supports both PSK and enterprise based authentication mechanisms for third-party authentication key management (e.g., LDAP directory).
- **Key Length:** WPA/WPA2 provide a minimum 256-bit pre-shared key length.
- **Message Integrity Code:** WPA correctly provides message integrity through a hashing mechanism knows as an MIC or Message Integrity Code. This is contrast to WEP's Integrity Check Value (ICV), which is comprised of an encrypted CRC-32 error-check. WEP's ICV is plagued by bit-flipping attacks due to its incorrect use as an integrity mechanism despite CRC-32's intended use as a transmission error detection mechanism
- **Mutual Authentication:** WPA/WPA2 provides the ability for an AP to be authenticated by a client through the use of certificates and RSN.

The core vulnerability which plagues WPA/WPA2 is not necessarily a vulnerability in the implementation, but an attribute of the necessary 4-way handshake process (EAPoL) for basic pre-shared key authentication. PSK is not designed for enterprise use (but nonetheless finds itself in the enterprise), which is why solutions exists within WPA/WPA2 to eliminate this vulnerability at the enterprise level but not at the SOHO level.

Pre-shared key authentication is commonly used in absence of a backend directory normally accessed via RADIUS and LDAP protocols. This type of implementation is common for solutions, which are not intended to be used by a wide variety of clients and therefore the cost of implementing such a solution is not effective. Pre-shared key authentication relies on an out-of-band key-management system. WPA-PSK and WPA2-PSK do not provide any key-management or tunneled encryption of the pre-shared key authentication handshake process. The password hash procedure used by WPA/WPA2 combines the stored passphrase with a salt of the destination AP's SSID and sends it over in plaintext.

Reliance on a single attribute for access control means that each client is indiscernible from the next and no amount of accountability can be afforded to each user outside of a DHCP MAC address log.

33.5.6.1 WPA/WPA2 Key Hash Capture and Cracking. The WPA/WPA2 pre-shared key authentication process is a series of handshakes used to eventually transfer a hash of the passphrase to the AP. This process can be captured through monitor mode and used to extract the hashed passphrase. Since this process only occurs when clients are initially authenticated to the AP, an attacker can either wait for a new client or deauthenticate an existing client by spoofing management frames as described earlier in the chapter.

Once a WPA hash is obtained an attacker can attempt to crack it using brute force. Due to the salting of the hash with the SSID of the AP, rainbow tables cannot be used to assist in the cracking process. The use of a salted hash combined with a lengthy hash process makes for slow cracking of WPA hashes on CPU-based systems. An exhaustive dictionary attack on a CPU-based system can typically process between 5,000–15,000 pairwise master keys per second.

GPU-based cracking systems are becoming more popular, as programming interfaces such as Nvidia's CUDA opened the door for developers to access many components of a GPU that would normally require advanced programming experience. GPUs can be leveraged to greatly increase the probability of successfully cracking a captured WPA handshake to the tune of 50,000+ pairwise master keys per second for an average set of CUDA-enabled GPU graphics cards.

Tools used to capture and crack a WPA/WPA2 pre-shared key:

- **Airmon-ng**[31]**:** Used to place the wireless adapter into monitor mode.
- **Airodump-ng:** Used to capture traffic using monitor mode.
- **Aireplay-ng-0:** Deauthenticates a client from the AP and forces them to re-authenticate and perform the pre-shared key handshake. This tool is necessary to allow Airodump-ng to capture the pre-shared key during the re-authentication process.

33.5.7 WPA/WPA2: Wi-Fi Protected Setup (WPS). Wi-Fi Protected Setup allows users with little understanding of wireless security configurations to easily configure a pre-shared key configuration without necessarily logging into the router. WPS describes the protocol and *personal identification number* (PIN) sequence used to negotiate a WPA/WPA2 pre-shared key without the need to manually configure it on the client and AP. WPS is only provided on SOHO-grade APs and not enterprise-grade hardware. This vulnerability may apply to the enterprise environment through its

existence as a nonmalicious rogue AP, which an internal user has securely configured but is vulnerable to WPS.

Two types of WPS exist:

- Push-Button-Connect (PBC)
- 8-digit PIN

A PIN may be provided by the AP for input into the client WPS agent or the client WPS agent may provide a PIN for input into AP WPS prompt. WPS authentication is based on 802.1X EAP and a vendor-specific implementation of EAP, which uses a Diffie-Hellman Key Exchange to prove possession of the first half of the PIN between the AP and client and eventually the second half.

Because WPS divides the PIN into two separate authentication steps, this process severely limits the number of combinations for each half of the PIN. An attacker can derive the first and the second part of the PIN through abusing an EAP message, "EAP-NACK." This message is issued by the AP on failing to provide the correct first or second portion of the PIN. Brute forcing each *separate* portion of the PIN requires only 10^4 tries at most, so there are only $10^4 + 10^4 = 20{,}000$ possible combinations instead of the original keyspace of 10^8 or 1,00,000,000.[32]

On average, WPS cracking tools can recover the WPA/WPA2 plaintext password in 4–10 hours.

Tools used to perform a WPS attack:

- **Airmon-ng:** To detect target AP and place adapter into monitor mode.
- **Reaver**[33]**:** Brute forces the WPS EAP process to recover the WPA/WPA2 key.
- **Wpscrack.py**[34]**:** Alternative tool to brute force the WPS EAP process to recover the WPA/WPA2 key.

33.5.8 MS-CHAP-v2—Divide and Conquer. EAP implementations that leverage MS-CHAP-v2 have already been discussed as vulnerable to dictionary attacks with the use of EAP-LEAP or EAP-PEAP-MSCHAP with poor supplicant configurations. Traditionally, the challenge-response hash-cracking process where a challenge-response hash is provided to a cracking program along with a dictionary was the only feasible approach to exploiting this vulnerability.

Thanks to research by Moxie Marlinspike (also the author of SSLStrip), this process was recently expedited through a new understanding of the seemingly painfully obvious and much overlooked functionality within the MS-CHAP-v2 protocol. During MS-CHAP-v2, the MD4 hash of a user's password is bitwise divided to construct three individual DES keys (seven bytes each). A typical penetration-testing tool such as *Asleep* will input the MD4 hash, break it into its three DES keys and use those keys to encrypt a given plaintext dictionary value for comparison to the original ciphertext hash.

The traditional cracking approach would result in a total keyspace of $2^{(56 + 56 + 56)}$ or about 10^{50}—a very large number that's infeasible for incremental brute-force cracking. This is exactly the reason why dictionaries were traditionally used as a "best guess" approach, because incremental (character-by-character) brute forcing is just not an option for any reasonable amount of time due to dictionaries' salting through the use of a challenge.

The use of three DES keys is prime for exploitation. An MD4 hash is only 16 bytes, whereas three seven-byte DES keys brings us to a total of 21 bytes. How is the difference made up? Padding, which reduces the effective keylength of the third DES key to two bytes. If each of the two seven-byte and singular two-byte keys are cracked individually, the effective keyspace is reduced to $2^{56} + 2^{56} + 2^{16}$, or about 10^{17}—a decline in keyspace by a factor of 10^{32}. The interesting part about the MS-CHAP-v2routine is that each DES key is used to encrypt the same plaintext; therefore, during the cracking process we can use the same key input for each of the two DES functions, effectively reducing the amount of DES functions for each iteration to one, leaving the only dual operation as the comparison of the DES output to each of the two previously determined ciphertexts. This is a final effective keyspace of 2^{56} (10^{16}), which in 1998 could be cracked in 4.5 days on average, and on the massively parallel architectures commercially available by the early 2010s, in about half a day on average.[35]

Tools used to crack MS-CHAP-v2:

- **Airmon-ng:** To capture MS-CHAP-v2 keying material during the challenge-response process.
- **Aireplay-ng-0:** Used to deauthenticate a connected MS-CHAP-v2 client.
- **Chapcrack.py:** Used to parse/extract the keying material from MS-CHAP-v2 handshake.
- **Super Computer:** A hardware-accelerated cracking machine designed to crack a single DES routine (available through cloud services).

33.6 MITIGATING CONTROLS. This section provides enterprise-grade technical mitigating controls that should be evaluated in every wireless solution regardless of the vendor-specific implementation. These controls respond to the fundamental wireless-security issues discussed earlier and provide a layered defense as part of secure enterprise design.

33.6.1 Improvements. Wired access controls have remained relatively steady with the inception of 802.1X port-based authentication, which is now standard in most modern switches. 802.1X port-based authentication is fundamentally the same authentication function baked into the WPA security standard and combined with the *Temporal Key Integrity Protocol* (TKIP). WPA2 builds on WPA by adding even stronger encryption, which requires dedicated hardware for AES encryption. The 802.1X wired standard is fully implemented in WPA and WPA2 but reliant on the configuration as to whether it is leveraged, such as *WPA-Personal*, which uses a pre-shared key, or *WPA-Enterprise*, which uses a central authentication repository.

The use of wireless for endpoint connectivity can provide security, which rivals traditional wired architecture through the ease of administration and management of a thin-client AP environment and network choke-points for monitoring in close proximity to the initiating clients at the outermost edge of the network.

33.6.1.1 RSN. Complete RSN support in WPA2 and 802.11i draft support in the WPA standard is essential to ensure wireless confidentiality and integrity. To deploy RSN, any AP which can perform 802.1X authentication will suffice. Support for 802.1X on an AP is usually just a matter of understanding the EAPoL protocol and accepting a RADIUS server as a configuration parameter. The AP does not need to be configured

for any particular EAP type, as this aspect of the network is completely transparent to the AP.

In a typical wireless environment, there are two fundamental management functions at play: user authentication management and key management for encryption, initial keying material, and key distribution. In a nonenterprise implementation, the user and key-management functions can be combined. The fact that a client possesses a correct pre-shared key authenticates them to the AP, and the bits which make up the pre-shared key may be used in the encryption process to derive a pairwise transient key for the encryption of the 802.11 payload.

33.6.1.2 Authentication and Access-Control Key-Management Considerations

- Solid authentication key management is necessary to deliver secure communication of keys in contrast to a pre-shared key environment where keys must be distributed through an out-of-band mechanism. A key can be in the form of a user/password combination, certificate, or other identifying material such an RSA SecurID number and PIN.
- Leverage an existing directory-based infrastructure used to authenticate users to computer systems on the wireless LAN or rollout an endpoint PKI infrastructure. The use of EAP and a backend user repository for authentication through the RADIUS protocol is the silver standard method of authentication. A PKI infrastructure, which provides client public/private keys to every wireless endpoint effectively eliminates the possibility for an attacker to capture a password hash challenge response.
- Strong EAP implementations include a TLS tunnel established between the endpoint client and backend authentication server (e.g., RADIUS) before any identifiable information is exchanged.

33.6.1.3 Encryption Key Management Considerations. As described in this chapter, encryption key management is the Achilles' heel of confidentiality. Modern encryption algorithms have matured to the point of perfection, as with AES, and are only as weak as their weakest link—that is, input. The best key-management functions are available only with the latest evolutions of wireless technologies. WPA2-Enterprise's AES/CCMP is the de facto standard and should be the minimum requirement for any new implementation. WEP, WPA, and WPA-2 present multiple methods of key management with improvement in each generation of 802.11 authentication/encryption schemes.

33.6.2 Authentication Server and Client Certificates. Certificate-based authentication between a client and backend authentication server is the *de facto* standard of wireless authentication. X.509 Certificates are used for the authentication of server to the client (RADIUS as a broker), client to the authentication server, or both.

To leverage endpoint-user and computer-based certificates for authentication in place of directory-based password authentication requires the implementation of a full *Public Key Infrastructure* (PKI). PKI is necessary to provide and regularly update/associate certificates with endpoint devices and sign them using an enterprise root Certificate Authority (CA) or outside root CA such as VeriSign.

For details of PKI, see Chapter 37 in this *Handbook*.

The following are three of the most common uses for AP or client certificates within modern enterprise wireless implementations:

- Authentication Server Certificates Only:
  - A barebones RSN EAP implementation should require at a minimum a signed authentication server certificate. This ensures to connecting clients that they are performing authentication with a trusted AP and authentication server. In absence of a server certificate, clients are susceptible to MITM attacks, which can be used to capture EAP challenge-response hashes from the client. The RSN protocol ensures that a client must authenticate the authentication server before authenticating itself to the server. This sequence protects against RADIUS-impersonation attacks and rogue APs if configured properly.
  - Wireless implementations which leverage authentication-server certificates but do not enforce them at the supplicant level forgo any security provided by this control. An attacker can masquerade as the legitimate backend authentication server through the use of same-SSID and hijack the authentication process.
- **Client and server mutual certificate-based authentication:** Each endpoint client is provided a certificate signed by the enterprise root authority and installed in the Operating System's certificate store. EAP-TLS is commonly used with this type of PKI infrastructure to authenticate the server's certificate (as with directory-based authentication) and the client's certificate by validating the certificate's signature. This client certificate also contains field information to identify the user and any other applicable information for use by the authentication server.
- **Client- and server-based authentication with external client certificates:** A PKI infrastructure can provide even greater security in the form of "something I have that others don't." A client certificate (private key) can be stored on a smart card, requiring the user to possess the certificate separate from the operating system in the event the laptop is lost or stolen. The smart card must be inserted into the computer prior to attempting authentication.

33.6.3 Endpoint Supplicant Configurations. The endpoint-client supplicant configuration is a crucial and commonly overlooked component of a secure wireless deployment. Deploying a *secure* wireless network within the enterprise protects the internal network, but most deployments neglect to assess the security of the endpoints that connect to the new wireless network.

Enabling wireless connectivity at the enterprise levels means that laptop computers, which are taken out of the office, will almost always have their wireless radios on. A wireless radio is not an inherent risk, but when combined with default operating-system and driver configurations, such unprotected devices can easily provide all kinds of connectivity information to listening parties.

Most default wireless-client supplicants lend themselves to some sort of attack, which can potentially compromise the endpoint at a remote location, such as a coffee shop, and re-establish connectivity when the client goes to work and connects the infected machine to the enterprise wireless network.

Employees may be flying, sitting on a bus, or just grabbing something to eat—and if they sit in a physical location long enough, an attacker could glean a wealth of information from an insecure supplicant configuration. This information can range from a pre-shared key hash for a home network, remembered SSIDs (e.g., the

enterprise wireless SSID), or even password hashes for enterprise directory-based authentication.

The following are recommendations when for securing an endpoint supplicant configuration:

- **Disable SSID Probe Requests:** Many operating systems and third-party supplicants will continuously probe for the last known SSIDs. This process is performed to ensure the client can connect to an AP, which may not be sending beacon frames to identify itself or if the client has missed the AP's beacon interval. A beacon interval is defined on the AP as the period of time between sending beacon frames to identify itself to all nonassociated clients within range. Disabling this feature prevents an attacker from passively monitoring 802.11 traffic and acquiring a list of SSIDs to which a particular client is attempting to connect. This information can then be leveraged to coerce the client into connecting to a rogue AP masquerading as one of the identified APs.[36]
- **Wireless Profiles:** Wireless profiles should be configured to separate corporate AP profiles from hotspot or home AP profiles. Third-party supplicants commonly provide this capability to prevent the situation where a corporate employee uses their computer at a coffee shop and their wireless driver unknowingly broadcasts the corporate SSID, therefore inviting an attacker to masquerade and potentially capture a WPA/WPA2 pre-shared key hash or challenge response directory-based password hash. The use of different profiles limits the ability for the client to know how to connect to APs on the corporate network when not located at the office.
- **Secure Authentication Server Configuration**
  - **Certificate Authority:** A supplicant should only accept authentication server certificates, which were signed from a designated/single root CA. This situation prevents the ability to connect to an AP and underlying authentication server, which is displaying a valid/matching Common Name signed by a legitimate but different root CA.
  - **Certificate Validation:** Users should not have the ability to overwrite and enable the use of an untrusted server certificate during the authentication process. By default, many supplicants will prompt the user if an invalid certificate is detected and provide the option to ignore. The supplicant should be configured to not prompt the user and remove any possibility for successful authentication attempts to an AP with an invalid or untrusted certificate.
  - **Common Name (CN):** Ensure the backend authentication server's common name or CN is specified in the client. This will prevent a connection to an AP and authentication server which possess a trusted certificate (assuming a nonenterprise-specific root CA) but uses an incorrect hostname/FQDN.

33.6.4 Wireless Intrusion-Detection Systems. The wireless intrusion-detection system (WIDS) has come a long way to become a standard network-security monitoring appliance, becoming an everyday security-monitoring tool, like a traditional wired IDS. A WIDS is capable of monitoring the 802.11 airwaves at the data link and MAC layer regardless of encryption or authentication. Although a WIDS cannot be used in place of a wired IDS, it can provide visibility into a cloudy segment of the network edge. A WIDS is also not exclusive to enterprises that deploy a wireless network. A WIDS can also provide value for enterprises, which don't

deploy wireless networks by ensuring against the unauthorized use of wireless connections by continuously monitoring the airwaves and alerting appropriately for internal rogue APs.

Multiple large breaches have stemmed from poor wireless management and visibility; this is a large problem for organizations with many physical retail or corporate outlets. Due to the widespread use and ease of installation of an SOHO AP obtained from a consumer retail outlet, it is becoming increasingly important to detect unauthorized extensions of the corporate network. In addition, the increased reliance on wireless communications for various business cases (such as corporate access, vendor devices requiring their own AP, or retail devices requiring dedicated APs with low-grade encryption due to low process power) has resulted in increased difficulty inventorying a network's authorized APs and confirming the proper segmentation or absence of internal connectivity for noncorporate access APs. In the event an enterprise can manage and inventory all known APs, there remains no native solution to quickly and accurately audit the implementation of enterprise standard authentication and encryption levels from a remote location.

33.6.4.1 Types of WIDS. Most modern WIDS infrastructures can be configured to leverage existing APs or use dedicated sensors to monitor 802.11 traffic. Dedicated sensors provide much more functionality at the expense of additional cost. A dedicated sensor has the time and CPU to solicit useful data from neighboring APs and clients, whereas an existing enterprise-thin AP may allocate only a fraction of processing toward these functions.

33.6.4.2 Use Cases. Due to the encryption, which takes place at the data link layer and above, a WIDS cannot inspect data-layer traffic. The majority of wireless threats involve the exploitation of a vulnerability, which exists in the data-link layer such as the 802.11 protocol itself or following authentication and encryption processes. The core value of a WIDS lies in the ability to provide assurance for wireless controls by verifying their existence and continuously monitoring for unauthorized changes.

A dedicated WIDS can provide the following services:

- Identify unauthorized APs across all physical locations (rogue APs).
- Detect and block the use of various wireless hacking tools with identifiable layer 1/2 signatures.
- Identify and continuously monitor for below enterprise-grade encryption or authentication mechanisms on authorized APs.
- Detect and block unauthorized connections to internal APs from clients with unauthorized wireless cards.
- Determine client supplicant configuration parameters as they relate to the transmitting of known SSID beacons.

33.6.4.3 Detailed WIDS Benefits

Regulatory Compliance. Regular wireless penetration testing and rogue AP audits are a necessary function of many regulatory certifications (such as the Payment Card Industry [PCI]). Leveraging a remote dedicated WIDS sensor may be a cost-effective alternative to deploying contractors or internal employees to each physical location to survey and assess controls.

Rogue-AP Detection. Rogue APs are the number-one issue threatening wireless-network security.[37] Even in the face of 802.1X switchport access controls, Port Address Translation (PAT) can provide access to the wired network for multiple wireless clients by sharing a single authorized MAC address. A single unmanaged AP could provide unlimited access to the internal network for an extended period.

Methods of rogue AP detection:

- Corporate AP white-listing by:
  - Vendor: Rogue APs can be detected based on a nonstandard make/model by performing a lookup on the base clients OUI portion of the MAC address.
  - SSID: Rogue APs can be detected through the identification of transmitted SSID beacons containing a nonstandard enterprise SSID (such as "Linksys" or "Netgear").
  - MAC Address: A static list of authorized AP MAC addresses can be obtained through SNMP or other means and provided to a WIDS system as "known good" APs.
  - Encryption/Authentication Grade: The WIDS can attempt to connect to potential rogue APs and enumerate available/requested encryption/authentication levels. A nonstandard level or method would be indicative of a rogue AP.
- Associated clients
  - A WIDS can perform analysis on a potential rogue AP's connected clients to determine if clients are enterprise-provided based on MAC address OUIs.
- Signal strength
  - A WIDS is constantly monitoring the airwaves and can alert on anomaly detection of a new AP in close proximity to the physical sensor/location by determining the signal strength of the potential rogue AP measured in dBm. If the dBm strength falls within a user configured threshold, an alert can be triggered.

Confirming internal LAN connectivity:

- Internal Device Connectivity: The remote WIDS sensor can attempt to connect to the rogue wireless network and ping a known internal device to confirm connectivity to the enterprise network (assumed no authentication).
- Simple Network Management Protocol (SNMP) content-addressable memory (CAM) table[38] enumeration: Internal network switch CAM tables possess all MAC addresses a device knows about and which port they reside on. This detection technique is based on the fact that wirelesses APs possess two network cards. One network card resides on the wired side of the AP, and the other on the wireless side/radio. These cards both possess MAC addresses, which are one or two values apart. A WIDS sensor can identify the MAC address in use by the wireless radio through passive monitoring of the data link layer and search all known switch CAM tables via SNMP for a similar MAC address. The WIDS controller can be provided an SNMP read string to the enterprise-wide switch environment used by the network management team for various monitoring/polling processes.

Counter-Action. If a rogue AP is detected, the WIDS can be configured to disable the switch port it resides on through SNMP write strings or to trap an alert to the central

console. A WIDS sensor can also be placed into a *blocking mode*, which will leverage spoofed 802.11 deauthentication packets to prevent clients from successfully establishing a connection to the rogue AP. Spoofed deauthentication packets were covered previously in Section 33.5.3, Management Frames, and in the correct circumstance can be used to *increase* security. These actions should be carefully analyzed prior to execution to minimize the probability of unknowingly attacking a legitimate neighboring business or household.

Defense in Depth—Rudimentary Controls. The previously described controls are best practice that can make or break the security of an enterprise wireless network. In addition to these controls, there are some controls which by themselves do not afford much protection, but together can reduce an attacker's success rate.

Security through obscurity may be a bad idea if it's the only security measure, but combined with solid security controls, one can potentially reduce legitimate attacks against core security control (such as EAP-MS-CHAP-v2). These controls are what we refer to as *script-kiddie* controls; they offer little resistance to an experienced attacker but cause an adversary to take multiple steps before even exposing their target objective.

- **MAC-address based restrictions:** Depending on the capabilities of each vendor-specific AP, it may be possible to restrict association to an AP based on MAC address, even in an enterprise environment. Methods exist to dynamically update MAC address mappings within a backend repository, similar to wired MAC-based VLAN assignment.
- **SSID Cloaking:** SSID cloaking is a perfect example of security through obscurity. The presence of an AP and BSSID can be identified through various penetration testing tools capable of analyzing raw 802.11 traffic, but an everyday operating-system default or third-party client supplicant will not display the presence of an AP with nonbroadcasting SSID.
- **Client Isolation:** Client isolation is an AP-specific control that monitors the destination MAC address of the incoming traffic. If traffic is destined for another client on the same AP, or SSID in thin client environment, it will be denied access. The purpose of this control is to limit the ability of a malicious user to communicate with other devices connected to the AP. This feature is common on public APs where the end user is not necessarily well vetted or authenticated. Depending on the enterprise use for a particular SSID/VLAN, this feature may be enabled to prevent access to other endpoints in an environment where services are provided over wireless by other wireless clients, such as a printer; analysis should be performed to allow only a subset of communication with those devices.

33.7 SECURE ENTERPRISE DESIGN. Earlier we discussed fundamental and specific vulnerabilities in detail to expose the underlying inherent risk of wireless network mediums and provide a background on some of the latest and greatest threats that should be addressed. Understanding an organization's risks and associated threats as they pertain to their unique network environment is vital to ensuring they are properly managed. The proper management of technical security risk starts with a secure design and ends with the implementation of secure technical controls to further support a secure design. The phrase "crunchy on the outside, soft and chewy on the inside" plays directly into this particular type of governance. Mitigating controls can

provide a crunchy perimeter, but secure design prevents a chewy interior in the event of perimeter-control failure.

In this section, we discuss priority design considerations which foster a secure enterprise wireless network. These design considerations should be accompanied by the technical mitigating controls discussed earlier to provide a wholesome approach. Integrating and allowing for various types of technical mitigating controls such as EAP methods, SSIDs, PKI infrastructures, AP certificates, or supplicant configurations during the initial wireless design phase can increase their effectiveness as they function as a whole in contrast to after-the-fact add-ons.

33.7.1 Benefits of Wireless Controller Architecture. Enterprise wireless network hardware comes in two primary varieties, thin (lightweight) and thick (autonomous) APs. These two types of designs will have a large impact on the way the wireless network is operated and managed, and the extent of available controls to secure it.

- A thick AP possesses a lot of intelligence and features outside of the core 802.11 or RSN functions. These features include helper services such as DHCP, SNMP, QOS, Firewalls, etc. Although these types of services definitely have a place, from a management standpoint it can become difficult to synchronize configurations across thick clients or varying vendors.
- Thin client APs are bare-bone APs, which implement the necessary 802.11/RSN functions only. Enterprises should leverage thin-client APs to lower cost and increase secure management of the wireless network. These APs leverage connected switch hardware or dedicated wireless controllers for intelligence and configuration settings. Upon boot and initialization a thin client will discover, contact, and download the latest configuration settings.

The security afforded by a thin-client environment reaches beyond ease of device management by providing the following additional benefits:

- **Physical safeguard of configuration data:** A thick client's configuration could be compromised through physical theft and analysis. A thin client holds the bare bone configuration to allow 802.11 connectivity, RADIUS shared secrets, and SNMP strings; other stored configuration information is held on the backend switch or wireless controller
- **Mesh Security Functions:** A common form of communication across APs such as the *Lightweight Access Point Protocol* (LWAPP)[39] enables widespread communication between APs and controllers across an entire enterprise. This communication mechanism combined with the physical presence of APs across primary and remote locations enables the ability to leverage them for security functions. Leveraging client APs for security monitoring can be beneficial from a cost-saving standpoint in favor of a dedicated WIDS, albeit with a lack of functionality.
  - Detect rogue APs using AP infrastructure rogue-discovery techniques. Vendor-specific solutions exist to leverage client-servicing APs to momentarily scan for rogue AP beacons, which are not part of the LWAPP network.
  - An AP can be placed into a detection role and connected to a switched port analyzer (SPAN) port to track down potential rogue AP beacons for internal network connectivity.[40]

33.7.2 Network Segmentation. A basic wireless network should be segmented in two ways. First, the wireless network should be segmented from the internal LAN network, and second, wireless networks should be segmented between wireless networks with different business uses. In addition, during the segmentation process consideration should be provided for WIDS sensor placement with the ability to physically span all areas of the wired network.

33.7.2.1 Wired versus Wireless Endpoint Segmentation. Wired networks commonly assign VLANs on a port or MAC address basis. In the case of port-based VLAN assignment, management relies heavily on proper access-layer switch configuration, trunk-port configuration for uplinks, and pruning the propagation of sensitive VLANs for different physical locations. Operational changes may dictate the updating of numerous switches for relatively minor moves. In the event of wired MAC address based VLAN assignment, MAC-authentication-aware switch hardware and backend RADIUS servers for MAC->VLAN mappings must be used. MAC-based VLAN assignment can improve management capability and oversight but poses a significant security risk as they are susceptible to MAC-address spoofing, which may allow an attacker access to sensitive VLANs by spoofing an authenticated device's MAC address.

In a wireless environment, VLANs can be assigned on a per-SSID or per-user-group basis, and administration efforts to provision a user to a VLAN can be greatly reduced or completely removed. A user is no longer confined to particular physical location to gain access to a required VLAN, as the necessary SSID may be available anywhere on campus. In the event of expansion, deploying a new thin AP in close proximately will deliver the necessary VLANs, which a client supplicant profile is configured for. Since VLAN assignment can be controlled through minor client supplicant configurations and backend RADIUS privileges, segmentation can be determined during the asset provisioning process and require relatively minor management thereafter.

33.7.2.2 Segmentation from the Wired LAN. Due to public accessibility of wireless airwave transmission (encrypted or not), segmentation within the existing LAN should be addressed to mitigate the risk of complete internal access due to a lapse of wireless security. A solid network security posture assumes a particularly risky segment may be compromised and is designed to mitigate the damage from such a disaster.

The value in segmenting the wireless LAN from the wired LAN thrives on the ability to hedge against the following events, which could lead to LAN compromise from a remote location, such as the corporate parking lot.

- Malicious use of compromised employee username/password–based authentication credentials through spear phishing emails or other theft.
- Misconfiguration of wireless security controls, such as a weak EAP implementation (EAP-LEAP or absence of AP certificate verification in EAP-PEAP-MS-CHAP-v2) lending itself to dictionary or RADIUS impersonation attacks.
- Rogue AP "hotspotting" of employee laptops to establish connectivity and compromise the endpoint through exploitation of services or HTTP traffic MITM. Once the endpoint is connected to the legitimate enterprise wireless an attacker could gain access to the LAN.

The need for segmentation of wireless traffic is due to the inherent lack of physical controls for 802.11 traffic and range of accessibility. An external attack at a DMZ or other Internet-accessible device could lead to the compromise of DMZ assets, but controls typically exist in the form of multileg or sandwich-style firewall environments which limit the exposure from the DMZ network and prevent direct LAN connectivity. This type of segmentation should be applied to the wireless environment for similar reasons, albeit different attack vectors.

33.7.2.3 VPN over Wireless. Depending on the purpose of the wireless network, users may require the same access on the wireless network that is afforded on the wired LAN. In favor of exposing the wireless network to the complete wired LAN a VPN solution may be used. This architecture would allow wireless network access to the Internet, just the VPN gateway, or other basic internal assets (such as an intranet site), which force the use of a VPN gateway for any other resources.

An additional strength provided by a VPN on top of wireless authentication is defense in depth. To gain internal LAN access from a wireless network, an attacker must compromise and gain access to the VPN in addition to any layer 2 wireless authentication/access control mechanisms. An attacker who gains access to the wireless medium would likely never suspect the use of a VPN and assume the state of limited internal access or mistake the current state for the complete LAN.

An assessment should be performed to understand the business case where wireless assets need complete access to the corporate LAN, and explicitly allow access on an IP/port basis. Fine-grain segmentation can be obtained by levering SSID-based VLAN assignments, explained in detail in the following sections, to assign restrictions on a per-VLAN basis with supporting ACLs at higher points within the network.

33.7.2.4 Inter-Wireless Network Segmentation. Inter-wireless segmentation enables different user types to connect to different SSIDs provided by the same physical AP and be afforded different network-based access privileges. The process of segmenting each SSID into its own network usually involves VLAN tagging at the AP level to distinguish layer-2 traffic by forcing inter-VLAN communication to flow over a firewall or upstream router with VLAN/IP-based ACLs. Backend directory systems such as LDAP can be leveraged by modern APs to provide a decision on the VLAN assignment a particular user will receive—regardless of the SSID on which they currently reside. Dynamic backend directory-based systems are not completely secure due to the reliance on MAC addresses as a unique identifying agent.

This type of segmentation can be enforced at the network layer, through aforementioned VLAN firewalling, or through VPN access lists if employee wireless access is forced through a VPN tunnel to communicate with LAN assets.

33.7.2.5 Benefits and Drawbacks of SSID/VLAN-Based Segmentation. Benefits of SSID/VLAN-based segmentation include:

- Differential treatment for wireless users by group or SSID; for example,
  - An SSID of “corporate-sales” and “corporate-hr” may provide varying levels of access. It would be best practice to disguise and cloak the SSID names to obfuscate their purpose to an outsider.
- Ease of network administration by removing the need to assign VLANs on a per-port basis. SSID-based VLAN assignment only requires endpoint configuration

to ensure the endpoint provides the credentials necessary to authenticate to a particular SSID.

A drawback of SSID-/VLAN-based segmentation is increased wireless-network overhead and reduced performance due to multiple SSID memory requirements.

Summary of Wireless Segmentation Controls

- LAN segmentation: protects internal LAN in the event of wireless compromise through VLAN and wireless IP–based firewall ACLs.
  - VPNs can be used over a wireless medium to allow users to access LAN assets.
- Intra-wireless network segmentation:
  - SSID based: The use of separate SSIDs for different business cases can easily be mapped to a VLAN and provide varying levels of segmentation.
  - Group based: RADIUS-based VLAN assignment for wireless user's current group privileges.

33.7.3 User Segmentation. A RADIUS server can make advanced authorization decisions for a valid username/password match. Leveraging designated user groups in the authorization process can provide fine-grained control over which users are able to use a wireless network segmented in large sweeps. If it's known that domain admins, power users, or general service accounts should never be able to authenticate over a wireless connection, they can be implicitly denied through nonmembership in a wireless-users group. Any user group that requires the ability to authenticate over wireless should be a member of this group.

A default RSN implementation that can and will authorize any user within the provided directory can become a victim of attacks on default, inactive, guest, or test accounts within the designated directory. Test accounts may be vulnerable to weak canonical passwords (e.g., "test" or "password") due to their creation during early stages of software or directory development or through other means that bypass corporate password policy. When a username/password authentication mechanism is in play, an attacker may attempt basic brute-force attacks to gain access to the network through standard test or guest accounts. Many organizations don't feel this is a serious problem, but when a directory contains thousands of accounts, and organizations don't audit on a regular basis, it's easy for things to slip through the cracks.

The use of a designated wireless group:

- Decreases the surface area of an attack, even if an adversary is able to obtain a list of user accounts and leverage them to brute force a wireless network through the use of basic passwords.
- Decreases the number of endpoints with saved or remembered network profiles that can be attacked and used to capture username/password challenge hashes in a hotspot or RADIUS impersonation-style attack.

33.8 SECURITY AUDITING TOOLS.[41] This section looks at the open source and commercial tools available to help audit a wireless environment and to understand the tools that war drivers have to attack a network. Exhibit 33.11 lists the main programs available, including older software still mentioned in discussions, and the platforms they run (or ran) under.

EXHIBIT 33.11 802.11 Security-Auditing Software

Tool	Description and Home Page	Operating System
Aircrack	A program that breaks WEP encryption, injects packets (Aireplay), and performs a dictionary attack against WPA PSK. http://www.aircrack-ng.org	Linux Windows
Auditor	An obsolete security collection that had all the required drivers and security tools to audit or hack wireless networks. Replaced by BackTrack (see below).	Linux
AirMagnet	A commercial wireless LAN analyzer that runs on laptop PCs and pocket PCs. http://www.flukenetworks.com/enterprise-network/wlan-security-and-analysis	Windows Pocket PC
Airsnort	Outdated tool. Broke WEP encryption using the RC4 Key Scheduling Vulnerability and Korek optimizations. Last update was 2005. http://sourceforge.net/projects/airsnort/	Linux
Airopeek NX	A commercial wireless LAN analyzer that ran on laptop and tablet PCs. Replaced by OmniPeek (see below).	Windows
BackTrack	A Linux security distribution and is the result of the merging of the Whax and Auditor security collections. http://www.backtrack-linux.org/ http://www.remote-exploit.org/articles/backtrack/index.html	Linux
CoWPAtty	A WPA PSK cracker last updated in 2008. http://www.willhackforsushi.com/Cowpatty.html	Linux
Ethereal	A program that understood the format of packets and dumps out their contents. Could capture packets as well as process a log file created by other programs. Replaced by Wireshark (see below) in 2006 when original programmer hired by a new company.	Linux Windows
Kismet	Wireless auditing tool still being actively developed and improved. Latest updates at time of writing were on 2012-12-06. http://www.kismetwireless.net/	Linux, FreeBSD, NetBSD, OpenBSD, Mac OS X
Linux-wlan	The Linux drivers for Intersil Prism 2, 2.5, and 3-based cards used by open source software such as Kismet and Wellenreiter. Last updated 2004. http://www.linux-wlan.com/linux-wlan/	GNU/Linux
Sniffer Global Analyzer	High-end software for use with portable (hardware) analyzer tools. http://www.netscout.com/products/enterprise/Pages/default.aspx	Hardware modules & client software
Netstumbler	Detects broadcasting 802.11 networks. Netstumbler is not suitable as a network auditing tool because access points can be configured to avoid detection. Last update (v0.4.0) was in 2004. http://stumbler.net/	Windows

(*continued*)

EXHIBIT 33.11 (*Continued*)

Tool	Description and Home Page	Operating System
OmniPeek	Commercial, sophisticated network-analysis tool from WildPackets with wireless capabilities. Replaced Airopeek. http://www.wildpackets.com/products/omnipeek_network_analyzer/wireless_network_analysis	Mac, Windows, hardware unit
Wellenreiter	Provides similar functionality to Netstumbler in that it detects access points. However, Wellenreiter has a major advantage over Netstumbler in that it detects access points that Netstumbler cannot see. At time of writing, the last update was 2012-07-17. http://sourceforge.net/projects/wellenreiter/	Linux
Windows wireless client	A built-in wireless subsystem that displays a list of available networks.	Windows
Wireshark	Highly regarded network protocol analyzer. Version 1.8.4 released 2012-11-28. http://www.wireshark.org/download.html	Windows, Mac

33.9 CONCLUDING REMARKS. The convenience and increasing bandwidth of wireless networks will inevitably result in continued growth of both technological improvements and implementation. Network and security managers must continue to monitor these developments and take appropriate action to secure all of their systems against constantly evolving threats to information security.

33.10 ABBREVIATIONS AND DEFINITIONS. Most of these abbreviations and definitions originate from the ANSI/IEEE 802.11-2012 standard.

Abbreviation/ Term	Meaning/Definition
802.11	The ANSI/IEEE standard 802.11, 1999 edition. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.
802.1X	IEEE Std 802.1X-2001 defines port-based network access controls. It provides the means to authenticate and authorize a network device before granting access to network resources.[42]
Access Controller	The CAPWAP term for a wireless switch that controls and manages Lightweight Access Points (LWAP).
Access point (AP)	A specialized wireless client that bridges other wireless clients to the physical wired network or distribution system.
Authenticator	An 802.1X component that authenticates a network device before allowing it to access network resources.
Authentication Server (AS)	An authentication server is an 802.1X component that performs the actual authentication of the supplicant on behalf of the authenticator.

Abbreviation/ Term	Meaning/Definition
BSS	Basic Service Set. A set of clients controlled by a single coordination function (an AP).
BSSID	Basic Service Set Identification. Due to the ability for a physical AP to provide multiple SSIDs on one interface, a BSSID reference to the MAC address associated with that interface to uniquely identify it as the provider of the SSIDs.
CAPWAP	Control and Provisioning of Wireless Access Points. An IETF initiative to establish standards enabling interoperability between different vendors APs and wireless switches.
CBC-MAC	Cipher Block Chain–Message Authentication Code. A message authentication algorithm based on a symmetric block cipher, where each block chains into the next data block.
CCMP	Counter mode with CBC-MAC protocol. A symmetric block cipher mode providing data privacy using counter mode, and data origin authentication using CBC-MAC. This is the protocol required for compliance with the WPA2 standard.
CRC	Cyclic redundancy check.
CRC-32	A 32 bit CRC that is used in the 802.11 standard.
Distribution System (DS)	A system used to interconnect a set of basic service sets (BSS) to create an extended service set (ESS).
DMZ	Demilitarized zone. This refers to a network segment that separates and protects a trusted network from an untrusted one.
DSSS	Direct sequence spread spectrum. DSSS is an 802.11 modulation technique that spreads signal power over a wide band of frequencies.
DTLS	Datagram Transport Layer Security. Essentially TLS over UDP.
EAP	Extensible authentication protocol defined in RFC 5247.[43]
EAP-PEAP	Protected EAP. A draft IETF standard by RSA, Microsoft, and CISCO for doing EAP authentication protected by TLS.[44]
EAP-TLS	EAP–Transport Layer Security defined in RFC 2716.[13]
EAP-TTLS	EAP—Tunneled TLS authentication protocol. A draft IETF standard that tunnels EAP authentication through TLS.[45]
EAPoL	EAP over LANs. The protocol defined in 802.1X for encapsulating EAP messages exchanged between a wireless client and the AP.
ESS	Extended Service Set. A set of one or more interconnected basic service sets (BSSs) and integrated LANs that appears as a single BSS to wireless clients.
Exclusive-Or (XOR)	An add-without-carry logical operation. XOR is used in cryptographic algorithms partly because it is reversible by repeating an operation. A XOR B XOR A will result in B; A XOR B XOR B will result in A.
FHSS	Frequency hopping spread spectrum. A transmission method that sends bursts of data over a number of frequencies, with the system hopping between frequencies in a defined way.

Abbreviation/ Term	Meaning/Definition
IBSS	Independent Basic Service Set. An ad hoc wireless network where clients communication directly with one another rather than via an AP.
ICV	Integrity check value. A CRC-32 value used by the WEP algorithm to detect changes in the ciphertext and to ensure that the packet has decrypted correctly.
IV	Initialization vector. A 24-bit value prepended to the WEP key, which is used as the seed for the RC4 stream cipher as part of the WEP algorithm.
KSA	The key scheduling algorithm of the RC4 stream cipher.
LAN	An IEEE 802 local area network.
LWAP	Lightweight Access Point. A type of AP that is essentially an 802.11 wireless radio that passes the received packets to a wireless switch for processing.
LWAPP	Lightweight Access Point Protocol. The base protocol chosen for the CAPWAP protocol and used by wireless switches to manage and communicate with access points.
MAC	Media access control. This is the data link layer in a wireless LAN that enables clients to share a common transmission medium.
Mb/s	Megabits per second.
MIC	Message integrity code. A cryptographic integrity code to detect changes in a message.
Michael	The message integrity code (MIC) for TKIP.
MITM	Man-in-the-middle. An MITM attack is one where an attacker can sit between two communicating parties and monitor what is being exchanged, without being noticed.
MPDU	MAC protocol data unit. The packets/frames exchanged between clients via the physical layer.
MSDU	MAC service data unit. Information that is to be delivered as a unit between MAC service APs (SAP). If an MSDU is too large to fit into one MPDU then it will be fragmented into multiple MPDUs and reassembled into one MSDU by the receiver.
OFDM	Orthogonal frequency division multiplexing. A modulation technique used in the 802.11a and 802.11g standards that works by dividing data into several pieces and simultaneously sending the pieces on many different subchannels. This mechanism enables throughput of up to 54Mb/s.
PAE	Port Access Entity. An 802.1X object that operates the authentication mechanism in the participants.
PKI	Public Key Infrastructure.
PMK	Pairwise master key. The PMK is a master session key derived from the overall master key. For 802.1X authentication mechanisms (e.g., EAP-TLS), the key is generated during authentication. For Pre-Shared Key authentication, the PMK is the pre-shared key (PSK).

Abbreviation/ Term	Meaning/Definition
Port	Network access port.
PRF	Pseudorandom function. An RSNA-defined algorithm that is part of the key-generation function.
PSK	Pre-shared key. A static cryptographic key that is distributed out-of-band to all clients in the wireless network.
PTK	Pairwise Transient Key. An RSNA cryptographic key that is the source of working temporal keys that protect wireless messages.
RADIUS	Remote authentication dial-in user service, defined in RFC 2865.[46]
RSN	Robust Security Network. The name for the new security system implemented in the RSN draft standard. An RSN is one that only allows the creation of RSN Associations.
RSNA	A Robust Security Network Association is where two clients have authenticated each other and associated using the 4-way handshake protocol.
RSN IE	Robust Security Network Information Element. The RSN IE is contained in Beacon and probe frames, and lists the supported authentication mechanisms and cipher suites.
Supplicant	The supplicant is the 802.1X term for the network device that wants to connect to the network.
"The Standard"	In this chapter, the term refers to the ANSI/IEEE standard 802.11-1999 edition.
TKIP	Temporal Key Integrity Protocol. TKIP is a suite of algorithms enhancing the WEP protocol to provide secure operations for legacy 802.11 equipment.
TLS	Transport Layer Security defined in RFC 2246.
TSN	Transient security network. A TSN can form associations with both RSNA capable and pre-RSNA clients.
UDP	User Datagram Protocol. An Internet protocol used to send short messages to other systems. UDP does not provide any reliability or ordering guarantees.
WEP	Wired Equivalent Privacy. The protocol defined in the 802.11-1999 standard to encrypt data on a wireless LAN.
Wireless LAN	In this chapter, the term "wireless LAN" refers to 802.11 wireless local area networks using a radio physical layer.
WLAN	Wireless local area network.
WTP	Wireless termination point. The CAPWAP term for an access point
WPA	Wi-Fi Protected Access. WPA was the first 802.11 wireless security standard from the Wi-Fi Alliance, using the TKIP security mechanism in the RSN standard.
WPA2	Wi-Fi Protected Access v2. WPA2 is the Wi-Fi Alliance's name for the AES-based CCMP wireless security mechanisms in the RSN standard.
XOR	Exclusive-OR.

33.11 FURTHER READING

Anjum, F., and P. Mouchtaris. *Security for Wireless Ad Hoc Networks*. Wiley, 2007.

Edney, J., and W. A. Arbaugh. *Real 802.11 Security: Wi-Fi Protected Access and 802.11i*. Addison-Wesley, 2003.

Coleman, David D., and David A. Westcott. *CWNA: Certified Wireless Network Administrator Official Study Guide: Exam PW0-105*, 3rd ed. Sybex, 2012.

Gast, M. *802.11 Wireless Networks: The Definitive Guide*, 2nd ed. O'Reilly & Associates, 2005.

Geier, Jim. *Designing and Deploying 802.11n Wireless Networks*. Cisco Press, 2010.

Makki, S., P. Reiher, K. Makki, N. Pissinou, and S. Makki, eds. *Mobile and Wireless Network Security and Privacy*. Springer, 2007.

Vacca, J. R. *Guide to Wireless Network Security*. Springer, 2006.

Wrightson, Tyler. *Wireless Network Security A Beginner's Guide.* McGraw-Hill Osborne Media, 2012.

Yang, Xiao, Xuemin Shen, and Ding-Zhu Du, eds. *Wireless Network Security.* Springer, 2007.

33.12 NOTES

1. IEEE LAN/MAN Standards Committee, "IEEE 802.11™: Wireless LANs," IEEE Standards Association, 2013, http://standards.ieee.org/about/get/802/802.11.html
2. David Chernicoff, "BYTE Guide To In-Flight Wi-Fi," *Information Week: Byte Reviews*, December 14, 2012, www.informationweek.com/byte/personal-tech/wireless/byte-guide-to-in-flight-wi-fi/240144399
3. IEEE LAN/MAN Standards Committee, "IEEE Std 802.11-2012 Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE Computer Society, March 29, 2012, http://standards.ieee.org/getieee802/download/802.11-2012.pdf
4. See also Stephen McQuerry, "Understanding WLAN Security," Chap. 3 in *Interconnecting Cisco Network Devices, Part 1* (ICND1): CCNA Exam 640-802 and ICND1 Exam 640-822, 528, Cisco Press, 2007;. an extract is available at www.ciscopress.com/articles/article.asp?p=1156068&seqNum=3
5. Microsoft, "Shared Key Authentication," MSDN—Microsoft Developer Network, April 8, 2010, http://msdn.microsoft.com/en-us/library/aa916565.aspx
6. Sheila Frankel, Bernard Eydt, Les Owens, and Karen Scarfone, "Establishing Wireless Robust Security Networks: A Guide to IEEE 802.11i (SP 800-97)," National Institute of Standards and Technology Special Publications, February 2007, http://csrc.nist.gov/publications/nistpubs/800-97/SP800-97.pdf
7. Eric Griffith, "WEP: Cracked in 60 Seconds," Wi-Fi Planet, April 9, 2007, www.wifiplanet.com/news/article.php/3670601
8. Margaret Rouse, "TKIP (Temporal Key Integrity Protocol)," SearchMobileComputing, March 2006, http://searchmobilecomputing.techtarget.com/definition/TKIP
9. Margaret Rouse, "CCMP (Counter Mode with Cipher Block Chaining Message Authentication Code Protocol)," SearchMobileComputing, June 2008, http://searchmobilecomputing.techtarget.com/definition/CCMP

10. IEEE LAN/MAN Standards Committee, "802.1X-2010—Port Based Network Access Control," IEEE 802.1 Working Group, February 5, 2010, http://standards.ieee.org/getieee802/download/802.1X-2010.pdf
11. B. Aboba, L. Blunk, J. Vollbrecht, J. Carlson, and H. Levkowetz, "Extensible Authentication Protocol (EAP): Request for Comments 3748," Internet Engineering Task Force, June 2004, http://tools.ietf.org/html/rfc3748
12. VOCAL Technologies, "EAPoL—Extensible Authentication Protocol over LAN," VOCALSecure Communication, 2012, www.vocal.com/secure-communication/eapol-extensible-authentication-protocol-over-lan
13. D. Simon, B. Aboba, and R. Hurst. " The EAP-TLS Authentication Protocol: Request for Comment 5216," Internet Engineering Task Force, March 2008, http://tools.ietf.org/html/rfc5216
14. P. Funk and S. Blake-Wilson, "Extensible Authentication Protocol Tunneled Transport Layer Security Authenticated Protocol Version 0 (EAP-TTLSv0): Request for Comments 5281," Internet Engineering Task Force, August 2008, http://tools.ietf.org/html/rfc5281
15. Matthew Gast, "TTLS and PEAP Comparison," Opus One, April 13, 2004, www.opus1.com/www/whitepapers/ttlsandpeap.pdf
16. Microsoft, "PEAP-MS-CHAP v2," Microsoft TechNet Windows Server, March 31, 2005, http://technet.microsoft.com/en-us/library/cc779326(v=ws.10).aspx
17. Moxie Marlinspike, "Divide and Conquer: Cracking MS-CHAPv2 with a 100% Success Rate," CloudCracker: Blog, July 29, 2012, https://www.cloudcracker.com/blog/2012/07/29/cracking-ms-chap-v2
18. Cisco, "PEAP/EAP-TLS Configuration Scenario," Configuration Guide for Cisco Secure ACS 4.1, n.d., accessed January 10, 2013, www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.1/configuration/guide/peap_tls.html
19. Krishna Sankar, Andrew Balinsky, Darrin Miller, and Sri Sundaralingam, "EAP Authentication Protocols for WLANs," Cisco Press, February 18, 2005, www.ciscopress.com/articles/article.asp?p=369223&seqNum=4
20. Sankar et al., "EAP Authentication Protocols for WLANS," p. 5, www.ciscopress.com/articles/article.asp?p = 369223&seqNum = 5
21. INTELLIGRAPHICS Device Drivers, "Introduction to IEEE 802.11," INTELLIGRAPHICS | Developer Resources | White Papers, 2012, www. intelligraphics.com/introduction-ieee-80211
22. IEEE LAN/MAN Standards Committee, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications—Amendment 2: Fast Basic Service Set (BSS) Transition," IEEE Standards Get IEEE 802 Program, July 15, 2008, http://standards.ieee.org/getieee802/download/802.11r-2008.pdf
23. Niels Ferguson, "Michael: An Improved MIC for 802.11 WEP | document IEEE 802.11-02/020r0," IEEE 802.11 Documents | Documents | Wiki, January 2002, https://mentor.ieee.org/802.11/dcn/02/11-02-0020-00-000i-michael-an-improved-mic-for-802-11-wep.doc
24. Jim Geier, "Understanding 802.11 Frame Types," Wi-Fi Planet | Tutorials. August 15, 2002, www.wi-fiplanet.com/tutorials/article.php/1447501
25. Aircrack-ng, Home Page v1.1, April 24, 2010, www.aircrack-ng.org

26. Paul Rubens, "Secure Your WLAN With Aircrack-ng," Wi-Fi Planet | Wireless Security, December 27, 2007, www.wi-fiplanet.com/tutorials/article.php/3718671/Secure-Your-WLAN-With-Aircrack-ng.htm
27. Dino A. Dai Zovi and Shane A. Macaulay, "Attacking Automatic Wireless Network Selection," Dino A. Dai Zovi's Website, March 18, 2005, www.theta44.org/karma/aawns.pdf
28. Lisa Phifer, "The Caffe Latte Attack: How It Works—and How To Block It," Wi-Fi Planet | Tutorials, December 12, 2007, www.wi-fiplanet.com/tutorials/article.php/10724_3716241_2
29. Martin Beck and Erik Tews, "Practical Attacks against WEP and WPA," aircrack-ng Website, November 8, 2008, http://dl.aircrack-ng.org/breakingwepandwpa.pdf
30. "mister_x," "Airmon-ng," Aircrack-ng Website, October 31, 2010, www.aircrack-ng.org/doku.php?id=airmon-ng
31. Beck and Tews, "Practical Attacks against WEP and WPA."
32. Stefan Viehböck, "Brute Forcing Wi-Fi Protected Setup: When Poor Design Meets Poor Implementation, v3," .braindump—RE and stuff, December 26, 2011, http://sviehb.files.wordpress.com/2011/12/viehboeck_wps.pdf
33. reaver-wps, "reaver-wps: Brute Force Attack against Wifi Protected Setup," Google Code Project Hosting, n.d., accessed January 10, 2013, http://code.google.com/p/reaver-wps
34. Admin, "How to Hack WPA WiFi Passwords by Cracking the WPS PIN," Live Technology Guide | Null Byte | Linux Security, 2013, http://livetechnoguide.com/how-to-hack-wpa-wifi-passwords-by-cracking-the-wps-pin (URL inactive).
35. Moxie Marlinspike, "Divide and Conquer."
36. Microsoft, "Description of the Wireless Client Update for Windows XP with Service Pack 2: Article ID 917021 Rev. 8," Microsoft Support, October 9, 2011, http://support.microsoft.com/kb/917021
37. Karen J. Bannan, "Rogue Mobile Devices Threaten Enterprise Security," DefenseSystems Website, 2011, http://defensesystems.com/microsites/2011/mobile-wireless/personal-devices-connect-corporate-infrastructures.aspx
38. T. S. Adams, "What Is a CAM Table?" wiseGeek Website, 2013, www.wisegeek.com/what-is-a-cam-table.htm
39. P. Calhoun et al., "Lightweight Access Point Protocol: Request for Comments 5412," Internet Engineering Task Force, February 2010, http://tools.ietf.org/html/rfc5412
40. Cisco, "Rogue AP Detection under Unified Wireless Networks—Document ID 70987," Cisco | Wireless, LAN (WLAN), September 25, 2007, www.cisco.com/en/US/tech/tk722/tk809/technologies_white_paper09186a0080722d8 c.shtml
41. This section is an updated extract of a table from Chapter 33 ("802.11 Wireless LAN Security") by Gary L. Tagg in the 5th edition of this *Handbook*.
42. IEEE LAN/MAN Standards Committee, "802.1X-2010—Port Based Network Access Control."
43. Aboba et al., "Extensible Authentication Protocol (EAP): Request for Comments 3748."
44. A. Palekar, Dan Simon, Glen Zorn, and S. Josefsson, "Protected EAP Protocol (PEAP)—Internet Draft," Internet Engineering Task Force, March 22, 2003, http://tools.ietf.org/html/draft-josefsson-pppext-eap-tls-eap-06

45. Funk and Blake-Wilson, “Extensible Authentication Protocol Tunneled Transport Layer Security Authenticated Protocol Version 0 (EAP-TTLSv0): Request for Comments 5281.”
46. C. Rigney, S. Willens, A. Rubens, and W. Simpson, “Remote Authentication Dial In User Service (RADIUS): Request for Comments 2865,” Internet Engineering Task Force, June 2000, http://tools.ietf.org/html/rfc2865

CHAPTER 34

SECURING VoIP

Christopher Dantos and John Mason

34.1 INTRODUCTION **34·1**

34.2 REGULATORY COMPLIANCE **34·2**

34.2.1 Enhanced 911 34·2
34.2.2 Communications Assistance for Law Enforcement Act (CALEA) 34·3
34.2.3 State Laws and Regulations 34·3
34.2.4 International Laws and Considerations 34·3

34.3 TECHNICAL ASPECTS OF VoIP SECURITY **34·3**

34.3.1 Protocol Basics 34·4
34.3.2 VoIP Threats 34·4

34.4 PROTECTING THE INFRASTRUCTURE **34·6**

34.4.1 Real-Time Antivirus Scanning 34·7
34.4.2 Application Layer Gateways and Firewalls 34·7
34.4.3 Logical Separation of Voice and Data 34·7
34.4.4 Quality of Service 34·7
34.4.5 Network Monitoring Tools 34·7
34.4.6 Device Authentication 34·8
34.4.7 User Authentication 34·8
34.4.8 Network Address Translation and NAT-Traversal 34·8

34.5 ENCRYPTION **34·9**

34.5.1 Secure SIP 34·9
34.5.2 Secure Real-Time Protocol 34·9
34.5.3 Session Border Control 34·9

34.6 CONCLUDING REMARKS **34·10**

34.7 FURTHER READING **34·10**

34.8 NOTES **34·11**

34.1 INTRODUCTION. Whether it is referred to as Voice over Internet Protocol (VoIP) or Internet Protocol Telephony (IPT), the digitization of voice messaging has had and will continue to have an impact on society. Voice messaging is part of a shift that some are calling the Unified Messaging System (UMS).[1] Instant messaging, Short Message System (SMS) text messaging to mobile phones, voice communications, video conferencing, email, file sharing, and network presence are already available in applications that are shared by both the home user and large corporations. Users are empowered by freeing our communications from geographically stationary limits. For example, users can decide to work from home and have their office telephones ring into their laptops. Colleagues can view each other's computer screens from anywhere in the world with Internet access.

Skype, for example, reported about 50 million concurrent online users at the time of writing in mid-January 2013—a phenomenal growth in participation in a UMS created in 2003 (about 590 percent per year over a decade).[2]

Aside from convenience and capability, there are also great financial justifications for VoIP/UMS systems. They replace formerly expensive equipment with common microcomputers, resulting in substantial savings. Some individuals and small offices are finding that they no longer need to buy landline phones: mobile phones and VoIP software with excellent wireless headsets or earpieces are even better than traditional landline units. An August 2012 report indicated that "New York state has seen a 55% drop in landline phones over the past decade. Some estimate that landlines will be completely obsolete by 2025" and argued that VoIP was "probably the most widely used landline alternative, especially among businesses."[3] By December 2012, U.S. Government reports indicated that more than half of U.S. residents did not have or use landlines in their homes.[4] Offices are turning to VoIP in the form of Hosted Private Branch Exchanges (PBX) to replace conventional landline PBXs.[5]

Just as there are security issues with our current operating systems and applications, the popularity of VoIP has already drawn the attention of criminals. Most disturbing of all is that many current exploits can be used against VoIP.

34.2 REGULATORY COMPLIANCE. Since VoIP affects key areas of an organization, such as systems, applications, privacy, networks, transmissions, and end users, VoIP must be included in all systematic risk analysis and compliance verification.

Among other U.S. laws, the Gramm-Leach-Bliley (GLB), Sarbanes-Oxley (SOX), and Health Insurance Portability and Accountability Act (HIPAA) are particularly significant in any evaluation of VoIP implementation.

For more information on GLB and SOX, see Chapter 64 in this *Handbook*. For more information about HIPAA, see Chapter 71.

34.2.1 Enhanced 911. Enhanced 911 (E911) mandates a location technology advanced by the Federal Communications Commission (FCC) that enables mobile or cellular phones to process 911 emergency calls and enable emergency services to locate the geographic position of the caller. According to the FCC's Website:

> The FCC has divided its wireless E9-1-1 program into two parts—Phase I and Phase II. Under Phase I, the FCC requires carriers, within six months of a valid request by a local Public Safety Answering Point (PSAP), to provide the PSAP with the telephone number of the originator of a wireless 9-1-1 call and the location of the cell site or base station transmitting the call.
>
> Under Phase II, the FCC requires wireless carriers, within six months of a valid request by a PSAP, to begin providing information that is more precise to PSAPs, specifically, the latitude and longitude of the caller. This information must meet FCC accuracy standards, generally to within 50 to 300 meters, depending on the type of technology used. The deployment of E9-1-1 requires the development of new technologies and upgrades to local 9-1-1 PSAPs, as well as coordination among public safety agencies, wireless carriers, technology vendors, equipment manufacturers, and local wireline carriers.

In a simple solution, a VoIP provider takes the stored location information, passes that through its own call center, and routes it to the local 911. In a more sophisticated solution, some gateways can map MAC addresses to locations, and the 911 calls then can be passed to a PSAP.

Currently, E911 is not required for companies that are using VoIP for internal purposes only. This law is empowered by the FCC, but the Department of Homeland

Security (DHS) and the Federal Bureau of Investigation (FBI) appear to want it for surveillance and location tracking.

34.2.2 Communications Assistance for Law Enforcement Act (CALEA).

Electronic surveillance consists of either the interception of call content (commonly referred to as wiretaps) or the interception of call-identifying information (commonly referred to as dialed-number extraction) through the use of pen registers and trap and trace devices.[6]

The standards for surveillance of VoIP networks have been a joint effort of industry and law enforcement for some time.

In 2007, the Packet Technologies and Systems Committee (PTSC) published the *Lawfully Authorized Electronic Surveillance (LAES) for Voice over Packet Technologies in Wireline Telecommunications Networks*, Version 2 (Revision of T1.678-2004), which is a surveillance standard for basic VoIP.[7]

A key part of the U.S. Federal Communications Assistance for Law Enforcement Act (CALEA) enforcement and debate is determining who is responsible for compliance. According to the law, "all entities engaged in the transmission or switching of wire or electronic communications as a common carrier for hire."[8] However, a question arises when an organization builds its own VoIP solution; that is, does the organization then become a *telecommunications carrier* subject to the law? It would seem so, according to CALEA:

> [A] person or entity engaged in providing wire or electronic communication switching or transmission service to the extent that the Commission finds that such service is a replacement for a substantial portion of the local telephone exchange service and that it is in the public interest to deem such a person or entity to be a telecommunications carrier for purposes of this title.[9]

However, the law specifically excludes "(i) persons or entities insofar as they are engaged in providing information services; and (ii) any class or category of telecommunications carriers that the Commission exempts by rule after consultation with the Attorney General." Given the ambiguity of this exclusionary language, an organization should seek advice from appropriate legal counsel prior to creating its own VoIP network.

34.2.3 State Laws and Regulations. All states have laws concerning surveillance; 31 states specifically address computers and 14 refer to cell phones. Because a detailed discussion of this diverse environment is beyond the scope of this text, an organization should seek appropriate legal counsel not only concerning the state it is domiciled in, but also the states of any branches, divisions, subsidiaries, or affiliates. The National Conference of State Legislators provides links to the applicable laws of each state and a summary of the coverage (e.g., cell phones, computers, video, photos, etc.).[10]

34.2.4 International Laws and Considerations. Similar to the diversity in the United States, many countries have laws concerning surveillance and intercepts. As such, appropriate legal guidance and advice should be sought before installing or using VoIP networks in international locations. In particular, standards such as ISO/IEC 27002:2005 and privacy laws may be applicable, and may affect VoIP implementation and usage.[11]

34.3 TECHNICAL ASPECTS OF VoIP SECURITY.

The next section is designed to provide a technical overview of VoIP and related security issues. It starts with an

introduction to the protocols used and then progresses to associated threats. Following that is a discussion of best practices and then encryption.

34.3.1 Protocol Basics

34.3.1.1 Audio Stream Protocols: RTP and UDP. Real-time Transport Protocol (RTP) is the packet-based communication protocol that provides the base for virtually all VoIP architectures. Part of RTP is timing and packet sequence information that can be used to reconstruct audio streams. Like TCP, User Datagram Protocol (UDP) is a layer 4 network communication protocol. Both of these protocols provide the basic packet addressing needed to get from one network address to another. When dealing with VoIP, packet delivery time is of the essence. Both TCP and UDP add overhead to the organization's network communication. This overhead results in greater time delays to the communication. UDP adds less overhead as it provides comparatively little packet error checking. The compromise is that UDP will help deliver packets quicker but will also result in more lost packets. It has been observed that a packet loss of 10 percent spread over a VoIP call may be virtually undetectable to most users. However, all of the features that can make VoIP secure also add latency. Added latency means more lost or discarded packets, which then results in an unpleasant or unacceptable user experience. In the end, a slow packet is a lost packet.

34.3.1.2 Signaling Protocols: SIP and H.323. Session Initiation Protocol (SIP) is a protocol that is used to establish interactive multimedia sessions between users. In addition to VoIP, it is also used for video conferencing and online gaming. SIP appears to be the most commonly accepted form of establishing VoIP calls. Secure SIP (SSIP) is discussed in Section 34.5.1. Like SIP, H.323 can be used for video conferencing or VoIP call setup. While SIP appears to be the standard for new installations. One may find H.323 in use at enterprise scale installations that have large investments in older, analog communication equipment.

Briefly, the basics of a VoIP call include these details. A VoIP client application, whether on a personal computer or a dedicated handset, uses SIP or H.323 to set up the call. This call setup is an exchange of control parameters that may include encryption and compression algorithms to be used. This is called "signaling." Once the call is set up, the VoIP client uses RTP to start packetizing the voice data. The RTP packets are incorporated into a UDP packet that adds addressing and sequencing information. The UDP packets are collected and sorted by sequence number at the receiving station. Some systems use a "jitter buffer" to assemble and store the packets. The endpoint VoIP client then reads the jitter buffer and turns the RTP packets back into voice.

34.3.2 VoIP Threats. Although not an exhaustive list, these hacks represent some of the vulnerabilities likely to be encountered when using VoIP:

- SPam over Internet Telephony (SPIT)
- Eavesdropping
- Theft of service
- Man-in-the-middle attacks

34.3.2.1 SPIT. Most users have become accustomed to finding 10 or 20 (or 100 or 200) spam email messages in an inbox on daily basis. The thought of receiving an equal number of voicemails on a daily basis leads to the unappetizing acronym

SPam over Internet Telephony (SPIT). Would-be SPIT spammers may have been both encouraged and then disheartened when, in 2004, Qovia,

> a company that sells enterprise tools for VoIP monitoring and management, … applied for a patent on technology to broadcast messages via VoIP—and another one for a method of blocking such broadcasts. The broadcast methodology only works on a pure VoIP network, while most of today's services are hybrids of IP and traditional telephone lines.

The author explains that Qovia realized that broadcasting VoIP messages could be useful for agencies such as Homeland Security but could also be abused by spammers. Therefore, Qovia pledged to "incorporate its SPIT-blocking technology in future releases of its security products, while enforcement of its patent on broadcasting, if granted, could be used to shut down VoIP spammers."[12]

Sam Rozenfeld, writing in March 2012, summarized some of the measures being used to interfere with SPIT.[13]

- VoIP providers can collaborate with business clients to use network firewalls to block inbound SPIT using blacklisting of sources known to the provider as large-scale SPIT senders.
- It may be possible for the VoIP provider to identify robocalls—automated large-scale SPIT floods—and block them.
- Authenticating the presence of a human being before allowing the caller to complete a call can block SPIT.
- Blacklisting specific individual numbers can help individual users suffering repeated messages from the same sources.
- Since some SPIT sources use area codes from which no legitimate traffic would normally occur, blocking entire area codes may work.
- Whitelisting might work for individuals who know exactly who should be using their phone number, but it is unlikely to be acceptable for businesses who must accept calls from potential customers.

34.3.2.2 Eavesdropping. In an unsecured VoIP environment, eavesdropping is reduced to a task that is quite straightforward. Earlier, it was mentioned that RTP is the de facto protocol for VoIP communication. RTP adds unique sequencing information to its packets. Because of this, one could collect a number of RTP packets and then assemble them in a consecutive order, as a receiving station would collect these packets in a jitter buffer.

Knowing this, the first step in eavesdropping is to obtain a packet-sniffing tool, such as Ethereal from www.wireshark.org. One may then use Ethereal to perform a packet capture, or one may obtain sample capture files from the Ethereal Website. If the user chooses to perform packet capture, care must be exercised not to violate any privacy laws applicable to the user's network. Once obtained, Ethereal is used to sort out any RTP packets. Exhibit 34.1 depicts a packet capture filtered to show RTP packets. Notice that the first column at the left contains a list of sequential numbers. These numbers indicate the packet placement within the capture. However, farther to the right is the actual conversation sequence number. Once the RTP packets are assembled by sequence number, they can be saved as an ".au" file, which can be played on most computers. As shown with this example, eavesdropping on unencrypted VoIP traffic is not a complicated or expensive process.

Filter: rtp.payload | Expression... | Clear | Apply

No.	Time	Source	Destination	Protocol	Info
624	1444.509099	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28590, Time=1240
625	1444.579046	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28591, Time=1400
626	1444.582579	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28592, Time=1560
627	1444.588245	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28593, Time=1720
628	1444.590352	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28594, Time=1880
629	1444.625165	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28595, Time=2040
630	1444.627060	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28596, Time=2200
631	1444.664688	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28597, Time=2360
632	1444.671724	192.168.1.2	212.242.33.36	RTP	Payload type=ITU-T G.711 PCMA, SSRC=932629361, Seq=28598, Time=2520

EXHIBIT 34.1 Packet Capture Showing RTP Packets

34.3.2.3 Theft of Service. One of the classic theft-of-service attacks that has occurred was reported in June 2006. In this case, a brute-force attack yielded special access codes allowing attackers entry into the provider network. Once in, the attackers obtained router passwords and login credentials; they then programmed the network transport devices to implicitly accept and route VoIP messages from the attackers' server. The attackers then sold VoIP access to public providers, who then sold it to the public. In the end, 15 exploited companies were left to pay the bill for the calls that were routed out of their network.[14]

An organization's own employees could exploit its VoIP infrastructure as these attackers did.

34.3.2.4 Man-in-the-Middle Attacks. As with any technology that digitizes communications, VoIP that sends data without encryption is inherently open to manipulation if an attacker can intercept traffic, alter its content, and send it on its way to the recipient in a classic man-in-the-middle (MITM) attack.[15] Once an attacker has control of VoIP traffic, the attacker can not only eavesdrop but can also:

- Initiate calls to a third party and impersonate a caller by sending data that appear to come from a legitimate phone belonging to someone else
- Deflect calls to the wrong destination
- Intercept traffic in real time and generate simulated voice content to create misleading impressions or cause operational errors

This last point warrants expansion. An eavesdropper could collect the digital patterns corresponding to words or phonemes generated by a particular user during normal conversation. Using simple programs, it would be possible to generate data streams corresponding to any spoken sequence including those words or phonemes in real time, allowing the MITM to feed the recipient (or even both sides of a conversation) with distorted or invented information and responses. The potential for mayhem is enormous, especially if the conversation involved, say, emergency response.

34.4 PROTECTING THE INFRASTRUCTURE.

This section focuses entirely on VoIP networks. For general information on infrastructure protection, see Chapters 22 and 23 in this *Handbook*.

34.4.1 Real-Time Antivirus Scanning. Protecting the VoIP infrastructure would appear to be a routine decision. Remember that as with VoIP, a slow packet is a lost packet. Many of the routine measures used to protect a typical server will introduce latency into the voice system, leading to jitter and sporadic communication. One of the first routine measures to be sacrificed is real-time antivirus protection of the VoIP server. Some vendors will suggest that real-time scanning of the organization's entire VoIP server is next to impossible, and unsupportable. This is an unacceptable myth. The VoIP server must be scanned in a fashion that is at least consistent with other production servers. Requests from system administrators asking to disable real-time scanning is a common first step in creating issues with the VoIP system.

34.4.2 Application Layer Gateways and Firewalls. An organization's VoIP infrastructure is much more than a series of systems that are digitizing voice and forwarding packets. Its servers will have real-time contact with email servers and central authentication systems using Remote Authentication Dial-in User Service (RADIUS) or perhaps Active Directory. It may also have database systems devoted to logging call information or even recording the calls themselves. An attacker who gains access to part of the organization's VoIP infrastructure may be able to access the most sensitive parts of its network. Consider the use of application layer gateways (ALGs) or SIP/VoIP-aware firewalls to segregate the VoIP systems.

34.4.3 Logical Separation of Voice and Data. It would be ideal to have a separate network for the organization's VoIP system. Bandwidth issues would be minimized and troubleshooting simplified. In most instances, it may not be possible to make a business case justifying the installation of a separate set of cables and network gear. The VoIP handset on the user's desk or the VoIP softphone in the user's PC will share the same wire as the workstation. The logical separation of voice and data begins with assigning the organization's VoIP devices to a network subnet separate from the data devices. This is initiated by a Dynamic Host Configuration Protocol (DHCP) request from the user's handset. Part of this request allows the DHCP servers to distribute addresses based on hardware identification parameters. For example, connecting a laptop to the organization's network will result in the assignment of an address that is on a subnet different from the Cisco handset that is plugged into the same connection. This logical separation of voice and data allows the organization's firewalls to protect its VoIP infrastructure by screening out protocols and requests that are not voice related.

34.4.4 Quality of Service. The term "quality of service" (QOS) refers to a set of configurable parameters that can be used to control and/or prioritize communication through the network. Again, with respect to VoIP, a slow packet is a lost packet. QOS can be used to prioritize the VoIP packets so they can be delivered in a timely fashion. Most vendors will provide a choice of default QOS configurations to be used according to the organization's needs. Some even provide VoIP firewalls that are capable of buffering VoIP messages and retransmitting packets. Regarding QOS, it is not just one parameter that can be turned on to make the VoIP installation work properly; it is a series of parameters that may need to be tuned to fit the organization's exact needs. For more technical detail, see IEEE standards 802.1p and 802.1q.

34.4.5 Network Monitoring Tools. Best practices require a dedicated security operation center (SOC) watching the organization's networks for attacks. Whether

it is a 7 × 24 SOC or just one network administrator who does everything, it is absolutely critical to provide the tools and training necessary to detect attacks and to troubleshoot performance issues. With VoIP, the staff will be facing attacks using a new set of attack vectors and protocols. If the network administration is outsourced, review of existing service-level agreements to guarantee that the provider is capable of supporting VoIP is strongly recommended.

34.4.6 Device Authentication. Device authentication can be accomplished in a variety of ways. The simplest is to store a list of device MAC addresses on the organization's VoIP server and to authenticate all SIP requests through that list of addresses. The standard is to deploy devices to desktops without configuration. Upon connection, a technician enters a setup utility that connects to the VoIP server and then downloads a preconfigured image.

34.4.7 User Authentication. The organization's finance people likely will demand some type of call tracking and usage so departments can be charged or reviewed appropriately. Users can also be placed into different groups that the organization can configure. This group-level access can be used to limit services, such as long distance or international calling. It is common for a VoIP infrastructure to have ties to a central authentication server, such as Lightweight Directory Access Protocol (LDAP) or Active Directory. This central authentication will ease functions such as forwarding voicemail to a personal computer or cell phone. There are two common problems to watch for:

1. The *authentication interval* will be set by user. The organization can configure a demand for authentication to be hours, days, or months; 24 hours is suggested. Some users complain, and demand that their central credentials be stored in the handset for months so they only need to log in infrequently; this practice is generally undesirable.
2. Most handsets will come with *default accounts and passwords*. These accounts must be disabled or, at the least, strong password discipline must be maintained over them.

34.4.8 Network Address Translation and NAT-Traversal. Network address translation (NAT) is a technique commonly used by firewalls and routers to allow multiple devices on an internal network to share one IP address on the Internet. A user's internal address should be known only to systems on that user's own network. When connecting to an external network, the organization's router or firewall forwards the user's communication to an external address but replaces the user's private address with its public address. When the communication is returned, the firewall routes the message back to the correct private address. Similarly, consider the broadband router in a home. The user may connect a series of systems to this device, each with a unique internal address distributed by the router, but the ISP sees it as only one address. These internal or private addresses are stored in what is commonly called a "translation table." At a higher level, the process of getting a packet through a NAT device is called NAT-Traversal (NAT-T). Understanding how NAT-T issues affect VoIP is vital to understanding the danger of sending a voice call to the Internet.

Section 34.3.1.2 outlined how SIP is commonly used to set up a call. Once the call is set up, the actual audio stream typically is relayed via RTP/UDP. This is where

the issues start. The firewall does not have a problem with passing SIP traffic back and forth to the Internet, as the internal address of the VoIP device is stored in the translation table. However, the SIP signaling has passed on the private address of the user's VoIP device. This means that a device outside of the local network is trying to send the RTP/UDP audio stream to a fictional IP address. Essentially, NAT prevents VoIP from functioning.

Several work-arounds are commonly used. Sometimes the NAT device can be configured to provide VoIP support. Sometimes VoIP devices can be configured to work over otherwise open ports, overloading a common protocol such as HTTP, unfortunately, often with unintended side effects. Or VoIP proxy servers can be used on either side of the NAT in order to facilitate the traversal. Each of these solutions opens up its own security concerns, which should be carefully addressed; these concerns include the consequences of external proxy servers or creating anomalous traffic over other protocols, as in the overloading example.

34.5 ENCRYPTION. Encryption plays a critical role in communications security. For a general introduction to encryption, see Chapter 7 in this *Handbook*; for more details of public key encryption, see Chapter 37.

34.5.1 Secure SIP. Transport Layer Security (TLS) was sponsored by the Internet Engineering Task Force (IETF) to secure and encrypt data communications crossing public networks. It is intended to replace Secure Sockets Layer (SSL) as a widely accepted form of securing data communication. This protocol consists of a "handshake" and a "record." TLS was designed to be application independent so developers could choose their own way of initiating a TLS session.

Secure SIP is a mechanism designed to send SIP signaling messages over an encrypted TLS channel. A SSIP session is initiated by a SIP client contacting a SIP proxy and requesting a TLS session. The proxy returns a certificate that the SIP client then authenticates. The client and proxy then exchange encryption keys for the session. If the call is destined for another network segment, the SIP proxy will contact that segment and negotiate a sequential TLS session, so the SIP message is protected by TLS the entire time.

34.5.2 Secure Real-Time Protocol. Secure Real-Time Protocol (SRTP) is an enhancement of RTP that provides encryption, authentication, and integrity to the VoIP audio stream. The Advanced Encryption Standard (AES) originally was a block cipher; SRTP incorporates AES into the data stream with an implementation that utilizes it as a stream cipher.

Encryption is good but it does not protect the user or organization against replay attacks. SRTP uses a Hashed Message Authentication Code (HMAC-SHA1) algorithm to provide authentication and integrity checks. The MAC is calculated using a cryptographic hashing function in conjunction with a private key. SRTP uses one of the five Secure Hash Algorithms (SHA) designed by the National Security Agency. All five of these algorithms are compliant with requirements set in the Federal Information Processing Standards (FIPS).

34.5.3 Session Border Control. To this point, a number of issues affecting a VoIP deployment have been identified. Session border control (SBC) is a set of services that address VoIP issues related to security, QOS, NAT traversal, and network

interoperability. SBC collects real-time bandwidth statistics that can be used to allocate the network resources necessary to maintain the QOS desired. SBC will also support a number of NAT-T algorithms that will allow calls to be routed to public networks while maintaining the anonymity of internal resources. At the same time, SBC can accommodate both SIP and H.323. This allows the signaling protocol translation necessary to connect both types of networks.

34.6 CONCLUDING REMARKS.

VoIP provides expanded functionality and lower costs for corporate users, but managers must integrate security considerations into the architecture and implementation of all such systems to prevent interception, deception, and denial-of-service attacks. In addition, technologists must monitor developments in this rapidly changing field to keep abreast of new attack methodologies and countermeasures.

34.7 FURTHER READING

Boyter, B. "Voice-over-IP Sniffing Attack," 2003, www.giac.org/certified_professionals/practicals/gcih/0442.php (URL inactive).

Collier, M., and D. Endler. *Hacking Exposed Unified Communications & VoIP Security Secrets & Solutions*, 2nd ed. McGraw-Hill Osborne Media, 2013.

Davidson, J., J. Peters, and B. Gracely. *Voice over IP Fundamentals*. Indianapolis, IN: Cisco Press, 2000.

Endler, D., and M. Collier. *Hacking Exposed—VoIP: Voice Over IP Security Secrets & Solutions*. New York: McGraw-Hill Osborne Media, 2007. See also the associated Website, www.hackingVoIP.com

Kuhn, D. R., T. J. Walsh, and S. Fries. "Security Considerations for VoIP Systems." NIST Special Publication 800-58, 2005, http://csrc.nist.gov/publications/nistpubs/800-58/SP800-58-final.pdf

Long, T. "Eavesdropping an IP Telephony Call." GIAC Security Essentials Certification Practical Assignment, 2002, www.sans.org/reading_room/whitepapers/telephone/318.php

Miller, W. *Wireless Mesh Network Security: An Overview*. Death and Blind Media, 2013

Molitor, A. "Deploying a Dynamic Voice over IP Firewall with IP Telephony Applications," 2000, http://cnscenter.future.co.kr/resource/rsc-center/vendor-wp/aravox/aravox_deploying_dynamic.pdf (URL inactive).

Molitor, A. "Securing VoIP Networks with Specific Techniques, Comprehensive Policies and VoIP-Capable Firewalls," 2000, http://cnscenter.future.co.kr/resource/rsc-center/vendor-wp/aravox/aravox_specifictechniques.pdf (URL inactive).

Porter, T., B. Baskin, L. Chaffin, M. Cross, J. Kanclirz, A. Rosela, C. Shim, and A. Zmolek. *Practical VoIP Security*. Rockland, MA: Syngress, 2006.

Thalhammer, J. "Security in VoIP Telephony Systems." Master's thesis, Institute for Applied Information Processing and Communications at the Graz University of Technology, Graz, Austria, 2002; www.iaik.tugraz.at/teaching/11_diplomarbeiten/archive/thalhammer.pdf (URL inactive).

Thermos, P., and A. Takanen. *Securing VoIP Networks: Threats, Vulnerabilities, and Countermeasures*. Boston, MA: Addison-Wesley, 2007.

Verma, P. K., and L. Wang. *Voice over IP Networks: Quality of Service, Pricing and Security*. Springer, 2013.

VOIP Security Alliance White Papers: www.voipsa.org/Resources/whitepapers.php

34.8 NOTES

1. J. M. S. Wams, "Unified Messaging Atop a Cloud of Micro-Objects," (doctoral dissertation, Vrije Universiteit Amsterdam, December 18, 2012), http://dare.ubvu.vu.nl/handle/1871/39352
2. Jean Mercier, "50 Million Concurrent Users Online!" *Skype Numerology,* January 21, 2013, http://skypenumerology.blogspot.com/2013/01/50-million-concurrent-users-online.html
3. Kate Harrison, "How To Break Up With Your Landline." *Forbes* Website, August 8, 2012. www.forbes.com/sites/kateharrison/2012/08/08/how-to-break-up-with-your-landline
4. Stacey Higginbotham, "Over Half of American Homes Don't Have or Use Their Landline." *GIGaom,* December 26, 2012, http://gigaom.com/2012/12/26/over-half-of-american-homes-dont-have-or-use-their-landline
5. Enhanced VoIP Communications Inc., "What is a Hosted PBX?" *Easy Office Phone* Website, 2013, www.easyofficephone.com/resources/hosted-pbx
6. Electronic Frontier Foundation, "CALEA FAQ: What are the different kinds of communications surveillance methods, and how are they relevant to the current debate regarding CALEA?" *Electronic Frontier Foundation: Defending Your Rights in the Digital World*, n.d. https://www.eff.org/pages/calea-faq#2
7. IHS, "ATIS Releases LAES Standard for Internet Access, Services—ATIS PP-1000013," *IHS News & Analysis*, March 30, 2007, www.ihs.com/news/atis-laes-internet.htm
8. Alliance for Telecommunications Industry Solutions. "Lawfully Authorized Electronic Surveillance (LAES) for Voice over Packet Technologies in Wireline Telecommunications Networks, Version 2," *ATIS Document Center*, May 2006, www.atis.org/docstore/product.aspx?id=22771
9. Cornell University Law School, "CALEA: 47 USC §1001—Definitions (Preliminary Version)," *Legal Information Institute* Website, January 4, 2012, www.law.cornell.edu/uscode/text/47/1001
10. National Conference of State Legislatures, "Electronic Surveillance Laws," *NCSL* Website, 2013, www.ncsl.org/issues-research/telecom/electronic-surveillance-laws.aspx
11. International Trade Administration, "Worldwide VoIP Regulatory and Market Information," *U. S. Department of Commerce*, November 3, 2009, http://web.ita.doc.gov/ITI/itiHome.nsf/f3e8a6b8413b3e5d85256cc40075c5df/cb2a434afea6790485256d020053fef0/$FILE/voip%20worldwide%202009.pdf (URL inactive).
12. Susan Kuchinskas, "Don't SPIT on VOIP," *Small Business Computing.com*, August 4, 2004, www.smallbusinesscomputing.com/news/article.php/3399011/Dont-SPIT-on-VOIP.htm
13. Sam Rozenfeld, "No more SPIT!" *Business VoIP/VoIP Security*, March 25, 2012, www.telephonyyourway.com/2012/03/25/no-more-spit
14. "VoIP Hacker Arrested on Fraud Charges," *Technology News Daily,* 2006, www.technologynewsdaily.com/node/3252 (URL inactive).
15. See, for example, P. Thermos, "Two Attacks against VoIP," *Security Focus,* 2006, www.securityfocus.com/infocus/1862/1 (URL inactive).

CHAPTER 35

SECURING P2P, IM, SMS, AND COLLABORATION TOOLS

Carl Ness

35.1 INTRODUCTION 35·1

35.2 GENERAL CONCEPTS AND DEFINITIONS 35·1
- 35.2.1 Peer to Peer 35·2
- 35.2.2 Instant Messaging 35·2
- 35.2.3 Short Message Service 35·3
- 35.2.4 Collaboration Tools 35·3

35.3 PEER-TO-PEER NETWORKS 35·4
- 35.3.1 Dangers to the Business 35·5
- 35.3.2 Prevention and Mitigation 35·7
- 35.3.3 Response 35·8
- 35.3.4 Case Study 35·9

35.4 SECURING INSTANT MESSAGING 35·9
- 35.4.1 Dangers to the Business 35·9
- 35.4.2 Prevention and Mitigation 35·11
- 35.4.3 Response 35·13
- 35.4.4 Safe Messaging 35·13

35.5 SECURING SMS 35·14
- 35.5.1 Dangers to the Business 35·14
- 35.5.2 Prevention and Mitigation 35·15
- 35.5.3 Reaction and Response 35·18

35.6 SECURING COLLABORATION TOOLS 35·18
- 35.6.1 Security versus Openness 35·18
- 35.6.2 Dangers of Collaboration Tools 35·19
- 35.6.3 Prevention and Mitigation 35·20
- 35.6.4 Reaction and Response 35·22

35.7 CONCLUSIONS 35·22

35.8 FURTHER READING 35·23

35.9 NOTES 35·23

35.1 INTRODUCTION. Peer-to-peer (P2P) communications, instant messaging (IM), short message service (SMS), and collaboration tools must be directly addressed in any comprehensive security plan. The dangers are real, as is the probability that at least one of these technologies is in use in almost every organization.

35.2 GENERAL CONCEPTS AND DEFINITIONS. This chapter is designed to present enough information and resources to aid in integrating the defense of each function into the organization's security plan. A list of resources is provided at the end of the chapter to aid in further research.

35.2.1 Peer to Peer. Peer-to-peer networking, also referred to as P2P, is not a new concept or technology. The term was contained in some of the original designs and proposals for the Internet as an efficient and logical way to exchange information from one resource, or peer, to another, on a large interconnected network. Today, the term is most associated with applications that transfer multimedia files across the Internet.

Peer-to-peer networks generally consist of different computers, or nodes, that communicate directly with each other, often with little, if any, need for a central computer to control the activity. Often utilizing an application with a client-server appearance, the two computers set up a direct connection between each other for file transfer. A central indexing computer may or may not be needed to help these computers "find" each other, to index and publish their contents, or to facilitate the connection. However, what is most important, the two computers must have a direct, logical connection to transfer the file or files. File transport may take place over a local network (LAN), a wide area network (WAN), a value-added network (VAN), or via the Internet.

Peer-to-peer technologies and applications were much more common in the early days of networking when it was not financially possible for many organizations to have expensive servers and complicated network topologies. This is especially true of personal computer networks that performed simple file sharing from computer to computer in a one-to-one model instead of today's much more common one-to-many server-to-client setup. However, there are legitimate uses for peer-to-peer technologies. One common example is the sharing of Linux distribution software images. Peer-to-peer sharing of these often large ISO disk images requires much fewer resources for the distributor, because there may be thousands of computers distributing the software among themselves, instead of every user trying to download the file from a single server.

35.2.2 Instant Messaging. Instant messaging, or IM, has become one of the most widely used communication mediums and is on pace to overtake email as the preferred technology to communicate with others. This tool allows users to communicate with each other in a real-time, synchronous, instantaneous fashion via computer, tablet, or mobile device. Today's IM applications are nowhere near the first generation of IM. The concept of communicating, or chatting, in real time made its appearance on multiuser computer systems, when users could initiate a text-based conversation with each other. The most common example of this type of communication was a host-based system such as a mainframe environment or UNIX system using programs like *talk* or *ytalk*. Initially, users may have been restricted to messaging each other when logged into the same machine; eventually users were able to communicate with each other, either via Internet Relay Chat (IRC) or via an early online service like America Online. The first widespread uses of IM were made possible with the popularity of the PC with a modem and were used mainly for brief, informal, personal conversations.

With time, IM has become a business tool and, in some organizations, a necessity, especially with telecommuters. The need to communicate with colleagues, salespeople, clients, customers, and the like has transformed a gimmick technology into ubiquity. With this change, it is necessary for security management to change and adapt accordingly. Users are able to send messages, files, real-time streaming video and audio, utilize collaborative online whiteboards, and share desktops, almost instantly. Essential to the organization or not, IM can become a dangerous medium for security breaches.

35.2.3 Short Message Service. Short message service, or more commonly SMS, is another previously minor technology that has become ubiquitous and a large part of everyday life for many people. Although some mobile phone standards and companies had different ideas for the uses of SMS, a common early use was to notify customers of information one way, from the mobile phone provider to the user. A popular example was alerting the user of a missed call or voicemail message. Many carriers never dreamed customers would actually be able to send text messages from one mobile phone to another, nor did the carriers think users would ever *want* to do such a thing. The name, short message service, also implied a limited amount of text a message could contain. Originally, users were limited to 160 characters or less.

SMS has morphed into something much larger. The commonality of mobile phones has pushed the original concept far beyond its original meaning and function. Today, two-way communication between mobile phone customers, often on different mobile phone carrier networks, between customers and mobile phone providers, and between customers and other information systems has become a way of life. Customers expect instant, always-on, reliable SMS services. Most mobile phones are capable of SMS text messages, taking and sending pictures, instant alerts, and a number of other services that utilize or expand on the original concept of short message service.

35.2.4 Collaboration Tools. People working together have created a need for even more technology to aid them in completing their tasks. There are many products in today's market to facilitate sharing, collaboration, and organization of data. As some information security professionals joke, "Computers and technology are generally safe and secure, until you let a human near them." Humans are inevitable when it comes to collaboration tools and systems. Many collaboration tools and systems are designed to aid workgroups that are physically far apart. Once a system has requirements that contain the words "open," "via the Internet," or "access from anywhere," information security professionals cringe. Securing collaboration tools can be difficult, especially when it comes to balancing functionality versus security. A handbook would not be complete if the cloud were omitted, which is precisely where most of the current collaboration tools reside, including their data. These tools are also creating a tug-of-war between users and security professionals—users want to use these tools; security pros worry about the consequences.

File-sharing services such as Dropbox and Google Docs have put effort into securing the data stored and shared by individuals and groups. For example, Dropbox answers the question "How secure is Dropbox?" as follows:

> We have a dedicated security team using the best tools and engineering practices available to build and maintain Dropbox, and you can rest assured that we've implemented multiple levels of security to protect and back up your files. You can also take advantage of two-step verification, a login authentication feature which you can enable to add another layer of security to your account.
>
> Other Dropbox users can't see your files in Dropbox unless you deliberately share links to files or share folders. Dropbox employees are prohibited from viewing the content of files you store in your account. Employees may access file metadata (e.g., file names and locations) when they have a legitimate reason, like providing technical support. Like most online services, we have a small number of employees who must be able to access user data for the reasons stated in our privacy policy (e.g., when legally required to do so). But that's the rare exception, not the rule. We have strict policy and technical access controls that prohibit employee access except in these rare circumstances. In addition, we employ a number of physical, technical, and heuristic security measures to protect user information from unauthorized access.[1]

Nonetheless, there are serious security concerns about cloud-based file-sharing tools and Dropbox in particular. At a Black Hat EU conference in 2013, the paper "DropSmack: How cloud synchronization services render your corporate firewall worthless" caught the eye of writer Michael Kassner. The paper summary included the following points:

- "… [C]loud-based synchronization solutions in general, and Dropbox in particular, can be used as a vector for delivering malware to an internal network."
- "… Dropbox synchronization service can be used as a Command and Control (C2) channel."
- "… [F]unctioning malware is able to use Dropbox to smuggle out data from exploited remote computers."

The paper's author, experienced penetration-tester Jacob Williams, warned that if a bad actor has any access to a secured Dropbox folder, it is possible to synchronize a remote-access Trojan he wrote called DropSmack with all the shared Dropbox folders. The tool would allow the infiltration of the entire corporate network. Williams also warned that access to a Dropbox folder by employees using their personal computers raises legal issues:

> Many general counsels are more than a little worried about the appearance of authorizing us to pen test what could end up being be home machines. That's becoming a sticky issue with pen-testers these days as people open spear phishing emails delivered to the corporate email addresses on machines that may be privately owned.[2]

Integrated collaboration tools have another danger, especially for inexperienced users who don't take daily backups: deleting one or more files (or all of them) in a shared Dropbox folder will propagate the deletion to all users of the shared folder. If any one of the users keeps a daily backup, the entire group of users can be protected against disaster; if none of them do, they may be in serious difficulty. In a related issue, any user who moves the files out of the Dropbox folder into a local folder will wipe the data from all the other users' Dropbox folders too.[3]

In a 2012 article, Matthew J. Schwartz urged corporate uses to pay attention to Dropbox use among their employees. His five recommendations (with more details in the original article) are:

1. Monitor Dropbox use
2. Compare cloud service security
3. Beware lackluster security cloud practices
4. Treat Dropbox as a public repository
5. Beware insider data exfiltration[4]

A free tool, Cloudfogger, automatically encrypts data on the client side when it is uploaded to any external collaboration tool using 256-bit AES encryption. The tool then automatically decrypts the data when it is downloaded by an authorized user.[5]

35.3 PEER-TO-PEER NETWORKS. One of the earliest mass applications of P2P was for free file-sharing of music through Napster, LLC. Despite difficulties over copyrights and a subsequent bankruptcy, Napster's technology, in substantially the same form, is still in widespread use. Practical applications have expanded beyond music

downloads into the business world, such as allowing small groups of users to share files without the interaction of a systems administrator and distribution of open source software. Likewise, it may be possible that employees are utilizing the organization's high-speed Internet connection to supplement their at-home movie collection via P2P downloads.

35.3.1 Dangers to the Business. Using P2P technology without proper care and controls, an organization may face serious consequences. There are many threats to an organization that does not properly control P2P networking, as for any other network configuration or protocol. Many problems are discussed in Chapters 21, 25, and 26 in this *Handbook*. However, this section contains several important issues that information security management should consider while performing risk analysis and policy implementation for P2P networking.

35.3.1.1 Abusing Company Resources and Illegal Content. Organizations must have an acceptable usage policy in place, one that limits what employees can do with the technology resources provided to them. The policy should clearly state the kinds of technology and applications that are prohibited or restricted in specific ways. In most cases, P2P technology used to download music or videos for personal use will violate the policy.

P2P technology is a specific danger to company technology resources because the inherent nature of P2P technology is to use every resource to the maximum extent possible. For example, a single P2P application, configured properly, will use every bit of bandwidth that is made available to it. This would include LAN bandwidth, WAN bandwidth, and Internet bandwidth. One of the most popular uses for P2P technology still remains the sharing of extremely large files, especially multimedia files, including full-length movies. These large files can take hours to download in full. This fact can have an extremely negative impact on an organization's network infrastructure—including expensive Internet bandwidth.

In practice, a single P2P application has been demonstrated to completely saturate a 12-megabit Internet connection, virtually denying, or severely limiting, access to all other computers.[6] In this case, these dangers to the business are common to many areas of information security management:

- Threat to **availability**. If an organization's resources, including network resources, are not available, the business cannot properly function.
- Threat to **integrity**. If the organization's resources are crippled or misused by employees utilizing P2P technology, data may suffer from a breakdown of integrity and usability.
- Threat to the **organization's image**. If the organization's information systems and infrastructure cannot be relied on because of interruptions from P2P abuses, there is a risk of financial or public image degradation. Some organizations are not able to overcome a substantial loss of image, credibility, or both.
- Threat from **litigation**. It is common to see illegal content being shared via P2P technology; illegal music and video sharing is often credited with having made P2P technology popular. An organization may suffer legal troubles, including copyright and intellectual property suits, if its resources are involved with the sharing of illegal materials. Some anti-piracy groups have become extremely aggressive in combating illegal sharing of copyrighted content.

35.3.1.2 Loss of Confidentiality. There are many ways an organization may suffer a loss of confidentiality from P2P technology. One common mistake is a misconfigured P2P application. The case study in Section 35.3.4 describes one situation. However, the dangers of a misconfigured P2P application are real—it is quite easy for data to be shared inadvertently. When users are in a hurry or do not understand what they are doing, a P2P application may allow for unauthorized access to information because its restrictions are too lax or missing altogether. A common mistake in a Microsoft Windows environment may be to share the entire "My Documents" folder when a user intended to share only photos. A misconfigured or maliciously altered P2P application may also become a conduit or access point for an attacker to enter an otherwise secure network environment.

Another, less well-known and often overlooked threat involves the amount of data a P2P application can reveal to unauthorized persons. For example, a P2P application may offer detailed information about its host, including:

- Operating system, version, and configuration
- Corporate network address scheme, host naming convention, DNS information
- Detail about the P2P application version or build (useful for attackers to exploit known vulnerabilities in a "buggy" release or version)
- Network routes, privileged access from a certain PC to sensitive networks within the organization (behind a firewall, etc.)
- Open network ports in the organization's firewall

Although many of these examples may seem rather benign by themselves, the P2P application may be revealing information that an attacker can use as part of a bigger attack. Chapter 19 of this *Handbook* details how small pieces of information can be gathered and used together in an information security breach. The nature and functionality of P2P applications leaks sensitive information that otherwise would not be revealed.

All P2P applications are not created equal; a P2P application may be different from the user's expectations. Can the P2P application actually be a reliable, malware-free, secure application—especially when the application is a free download from the Internet? It is possible that a backdoor exploit, malware, spyware, or the like may be built into the P2P application, or introduced later. This was especially true in the days of Napster; many applications included unwanted malware that ranged from innocent to downright dangerous.[7] Similar exploits are still possible.

Many users still do not completely understand the application's functionality. What may start out as a user trying to download a single movie can turn into an application that never gets uninstalled and is always running, uploading, or "seeding" that single movie for weeks or months on end. Since many of these applications are built to silently operate in the background, the user may think that the movie is downloaded, and the application isn't used any longer, when in reality, it is still running until it is removed. This is not only a waste of resources, but a likely way to expose the organization to a DMCA (Digital Millennium Copyright Act) complaint from the organization that owns the intellectual property.

35.3.1.3 Consequences. Any organization that does not protect against data loss via P2P networking is at great risk of public disclosure and scrutiny, financial

penalties, regulatory penalties, and so on. The functionality and nature of P2P applications may provide an investigator or, worse yet, the press with definitive evidence of the use of P2P technology within an organization. (It is easy to discover an organization that has users utilizing P2P applications because there are several online databases of IP addresses recorded while participating in a P2P "hive.") A majority of the public may only understand P2P technologies to be used in conjunction with illegal music sharing; even this simple, negative perception can greatly influence public opinion on the organization. It would be difficult to refute packet analysis or screen shots containing an organization's IP address in which the computer was compromised, used for illegal software or media sharing, or the computer was used by an unauthorized entity to extract data. In the age of P2P applications commonly used to illegally share and distribute the intellectual property of unwilling participants, organizations are taking aggressive steps to find and prosecute offenders. See Chapter 55 in this *Handbook* for a discussion of cyber investigations; see Chapter 61 for guidance on working with law enforcement.

35.3.2 Prevention and Mitigation. Protecting the organization from information security breaches via P2P technologies is one of many important parts to an overall security plan. Depending on an organization's structure, leadership, function, and similar factors, methods for preventing and mitigating P2P threats can range from simple to complicated. Obviously, each organization must perform a risk analysis and determine its threat threshold when it comes to P2P technology. Chapter 62 provides means for risk assessment. The guidelines that follow can help an organization defend against the threat of P2P technology causing security breaches.

35.3.2.1 Policy. It is important for every organization to address the use of P2P technology in a policy, such as an acceptable use policy, HR policies, or security policies. The relevant policy, along with all other security-related policies, should be clearly stated, clearly communicated to the entire organization, uniformly and equally enforced, and updated as necessary.

35.3.2.2 Complete Ban on Peer-to-Peer Technology. In *most* cases, the organization can ban the use of P2P completely, especially through enforceable policy. Care should be taken to ensure all employees and computers are in compliance with the ban. It should be forbidden or, even better, impossible to install P2P applications on personal computers, servers, and all other information systems that could be used to send and receive P2P-related traffic. Most computer users should be running as a standard user of their computer, not as an administrator. If employees are allowed to install software, or are allowed administrative access to their desktop, regular, automated inventories and audits of the computers should take place. Removal should be immediate and appropriate corrective actions taken.

Several technologies may also aid in disallowing P2P traffic, although no technological solution is completely foolproof. These measures are additional safeguards, not complete solutions. Firewalls should be configured with a default deny policy and should only allow TCP/IP ports to pass that are necessary for normal business operations. While many P2P applications are able to tunnel through TCP/IP ports such as those used by HTTP or other common protocols, this is a necessary first defense. Packet-shaping technologies can also be useful to identify P2P-related traffic and block its communications. Packet-shaping and traffic management devices are often able to detect the signature of P2P traffic, no matter what TCP/IP port the application may

be using. Some intrusion-detection and intrusion-prevention systems may also be able to identify and block P2P traffic, as would many Internet filtering devices. Logs and reports should be examined daily and infractions should be quickly remedied. This is also a case where a managed desktop environment is important—automated software install/removal, uniform desktop imaging, and automated software inventory/reporting aids the desktop administrators in the effort of combating P2P software installed against policies.

35.3.2.3 Information Security and Information System Audits. All information systems and components should be audited (both on a schedule and via random audits) to ensure that they are not configured, intentionally or unintentionally, to participate in P2P file sharing. This task should be part of every organization's regular information system and information security audit processes. If possible, external and neutral resources are most useful to ensure all systems are audited in a uniform, exact, repeatable, and objective fashion.

35.3.2.4 Legitimate Business Use Must Be Managed. There are times when an organization does not wish to completely ban or block the use of P2P technologies. One increasingly common and legitimate example for P2P involves open source software distribution or updates via BitTorrent. BitTorrent is a P2P-based protocol for the distribution of data—often large amounts of data. A widespread use includes the distribution of several distributions of the Linux operating system. Installation of Linux often involves obtaining CD-ROM or DVD images to create install discs. By utilizing BitTorrent technology, software vendors and distributors are able to provide large amounts of data to their clients without carrying the entire burden of distribution, bandwidth, and computing resources. However, the organization must manage how this technology is used to ensure resources are not abused and that the P2P applications are used for only allowed, legal ends.

This can be accomplished through policy, auditing, and various network access and control technologies. Each organization must define its own level of acceptable risk for legitimate P2P technology usage and must find solutions that will match the acceptable level. Some examples include:

- Use of encryption
- "Anonymous" P2P routing technologies such an onion routing (see Chapter 31 in this *Handbook*)
- Network isolation for computers used to obtain software with P2P applications
- Company-acquired DSL or cable modem connections to the Internet, avoiding the use of corporate network resources

35.3.3 Response. It is necessary for all organizations to define exactly how to respond to security breaches and policy violations, including situations where P2P technology is involved. Not only should the process be included in the overall security plan, but also incident response processes should be in place to remove offending systems from the network. In some cases, the rebuilding of a compromised resource may be necessary, but some organizations may choose to remove the compromised machine from production and/or preserve a forensic copy of the machine for legal, forensic, or investigative processes.

35.3.4 Case Study. Misconfiguration, unintentional use, curiosity, and experimentation with P2P in the workplace do happen, with consequences. Although this case is only one type of specific security incident involving P2P technology, it should serve as an example of how such a situation can occur.

An employee of one organization reported a slow-running computer to the helpdesk. All of the usual helpdesk suggestions and tricks were exhausted with little effect on the performance of the computer. The usual symptoms of a slow computer were present—massive wait times to accomplish simple tasks, random errors and shutdowns, lockups, and other operational problems. However, there was one difference: After a reboot, it would take several minutes for the computer to slow down and become unresponsive. After some time, an employee commented, "I did try installing a music sharing program last week, but I didn't like it and uninstalled it." This led the engineer to examine each and every process that was running on the computer.

Although it appeared that the P2P application had been uninstalled, it actually had not been; it was still installed and running in a stealth mode. The uninstaller only masked the P2P application. Not only was the P2P application still running, but it was misconfigured to share the entire contents of the C: drive. There were literally thousands of other P2P users attached to the machine actively searching, uploading, downloading, and altering the contents of the computer's hard drive. The computer was not only giving away all of its data, it was being used for a server to host thousands of media files. Since the computer was on a network segment that had full TCP/IP 1-to-1 network address translation, it was effectively completely open to the outside world—and the outside world was taking full advantage of the opportunity. The hard drive was virtually full, and files were being added and deleted at will by remote users. The host's firewall was even modified by the P2P application's install to open all necessary ports to the world.

It is unknown if the user's personal data was actually accessed, downloaded, or used for any malicious activity, but the capability was certainly there. Because of a user's unauthorized download, inadequate network security, and other policy violations, the organization could not be sure of the confidentiality or integrity of the computer or its data. Necessary steps were taken to prevent this incident from occurring again, but this scenario has played out at other organizations, and will continue to do so as long as the P2P risk exists.

35.4 SECURING INSTANT MESSAGING.

Instant messaging has become an integral part of communications—both business and personal—for many people. It was fairly easy in the recent past to simply create a policy that banned IM for personal use and only allow internal, business-related IM within the company. However, IM has become an integral part of life. IM is pervasive in most organizations—for personal and business use. As a result from executives to interns, IM may be found on many desktops, but it must be managed and secured on all.

35.4.1 Dangers to the Business. With any technology, especially those that make connections to the Internet, there is a risk to the organization. Instant messaging is not a petty annoyance that should be taken lightly; if the technology is not controlled by the organization, a serious breach of security could occur. Remember, IM applications today go well beyond just txt; they are capable of transmitting much more than just interpersonal casual banter.

35.4.1.1 Loss of Information. Information loss and loss of confidentiality, intentionally or unintentionally, is most likely the biggest threat of IM to the business. There are several ways information can be harmfully conveyed via IM[8]:

- Revealing secrets via text chat, especially when a company has taken extensive effort to filter and block such communication via email with email gateways, intellectual property data management, and other technologies
- Copy-and-paste functions (including screen captures) used to transmit confidential or secret information from other secured applications or environments
- File transfers
- Screen sharing and real-time collaboration functions such as shared whiteboards or desktop sharing functions
- Forgetting to turn off a voice or video session and unintentionally relaying sound or images to the original correspondent—especially dangerous when leaving voice mail through the IM client
- Relaying voice, video, or both to another party (unintentional or intentional)
- Use of Webcam technology to relay visual information within a secure facility
- Downloading malware to collect and steal data
- Impersonation (This tactic usually involves stealing a known IM account or creating a fake account to impersonate someone the victim knows.)
- Subpoenas or search warrants executed to collect IM logs, conversations, and so on

Although this is not a complete list, it should serve to aid in security planning and policy development. There are many good resources on the Internet to further explain similar threats and consequences, but the preceding list should encourage thoughtful brainstorming about the ways in which an organization may lose data. Some of the listed methods would be extremely difficult to detect and remedy. With high-speed networks and high-speed Internet links at most organizations, a massive amount of data can be conveyed within a small amount of time.

35.4.1.2 Consequences. Like other security threats, the consequences of not securing IM technology can be serious. Many organizations today have experienced a security breach that involved IM, and there are probably more to come. Instant messaging security breaches can be deadly to an organization by themselves or as part of a much larger attack on the business. Stealing or transmitting information through IM is no less risky than any other form of information theft. One single file, whether sent through an IM file transfer or meticulously cut and pasted, bit by bit over a great period of time, can destroy a company's reputation and standing in the public eye or even benefit a competitor. A breach from a single IM conversation has the potential to depress a corporation's stock price in a matter of hours or days. There have even been cases where a chief executive officer's confidential information was captured and posted on the Internet for all to see.[9]

35.4.1.3 Denial of Service. IM cannot be written off as a tiny application with no real footprint on network resources. Instant messaging can be a tool used to create a denial of service attack on an organization, resulting in a loss of availability. IM

technology can be a powerful and useful tool for an attacker, including the use of IM clients with a direct connection to the Internet. With the right combination of malware and access, an attacker may be able to exploit one of many vulnerabilities discovered in IM applications, including the ever-popular buffer overflow. The National Vulnerability Database listed 38 current vulnerabilities in instant messaging software as of March 2012.[10]

35.4.2 Prevention and Mitigation. Every organization must guard against the threats caused by IM technology. Proper review and analysis of the risks associated with IM must be carried out within the organization, and the organization must determine the amount of risk it is willing to take. It is also necessary to evaluate the costs and efforts associated with the prevention and mitigation of this threat. Organizations will judge the risks and rewards of using IM differently. There is no set standard for every organization or business; there is no universal set of rules that can be applied in all situations. (For a discussion of risk assessment and management, see Chapter 62 in this *Handbook.*) The next sections provide strategies, tactics, and considerations for securing IM.

35.4.2.1 Policy. Policy must come first, especially with the popularity and widespread use of IM. Without adequate policies, the organization has no chance of actually protecting itself. Policy must be the foundation that all other considerations rest on. Clearly defined, well-communicated, and equally enforced policy is one of the most important fundamentals information security relies on. No matter what the organization decides when it comes to IM rules, it must be stated in a policy.

Instant messaging, while risky, is one of the most visible policy decisions a business will make for employees. While it might be best, and preferred for best security, to completely disallow IM, which could lead to frustrated employees, unable to use the facility for personal use, business use, or both. Every management team should be conscious of the potential ramifications of an overly strict policy. Conversely, allowing unfettered IM is certainly not the best solution.

Some organizations, depending on their software platform, may be able to provide employees with a managed IM environment where the employee can connect to the organization's official IM platform, as well as some of the popular IM services used for personal communications. At the same time, the organization can centrally deploy settings such as which services are allowed, how they connect (encrypted or not), and what logging options are in use—including if conversations can be recorded or archived.

Compliance and governmental regulations must be taken into consideration. If IM communications are to be allowed, they may become part of the organization's business communications, and therefore may be subject to subpoena, open records requests, and archiving and document retention requirements. This can be especially dangerous if the chosen IM platform is integrated into the organization's voice over IP system. New regulations and legal rules may greatly affect policy decisions. It is important to remember that instant messaging logs and conversations *may* be subject to legal discovery, search, and seizure. Consult counsel for proper legal advice.

35.4.2.2 Effects of a Complete Ban. A ban on all IM technology would be the best way to ensure better enterprise security. However, this will only produce dissatisfied users, without being effective. Users can become technology-savvy in a hurry if they are determined to circumvent a policy. A block on IM communications often

causes users to do just about anything they can to accomplish their goal of unobstructed IM. Some clients are designed with this in mind and will happily tunnel out of commonly allowed TCP/IP ports such as 80, usually open for HTTP communications. The software may be configured to bypass firewall rules, detect and avert packet-shaping technology, and tunnel its way to the Internet via encrypted connections. Many IM services also provide Web-only interfaces that do not require software to be installed while communicating via HTTP. Unfortunately, this technology can be a difficult and frustrating one to ban within the organization; a complete ban is probably not a practical solution.

35.4.2.3 Prevent Installation of Instant Messaging Software. Although a complete ban may not be possible, or even desirable, one step that a more secure organization can take is to prevent users from installing IM software. This tactic is not going to solve the whole problem, but it certainly will help. Controlling the installation of IM software should be part of the organization's overall software installation policy. In general, installing software without permission should be denied. If it is feasible, local workstation administrator rights should be denied for most employees. In many user or system management solutions, it is also possible to block software installation through individual, group, or workstation policy templates and procedures. Universal software installation prevention is much easier than trying to define policies or templates for every possible IM peer, application, group, or tool.

35.4.2.4 Fight Technology with Technology. Technology on its own will not provide the organization with an all-in-one solution for securing IM. However, there are a number of network devices, appliances, traffic monitoring software, and other technologies to help an organization minimize IM use. Do not believe marketing claims that any device or technology can guarantee IM blocking; few can deliver on this promise. The only *true* way to guarantee an IM-free company is to block access to the Internet completely—which is not realistic. An organization can also leverage existing security infrastructure such as IDS or IPS.

35.4.2.5 Limit Risk and Exposure. For most organizations, *limiting* IM through policy and technology is the solution to the threats that IM introduces. Combining those two approaches will help to reduce the possibility of data loss. Security managers and administrators should agree on what can and what cannot be allowed within the organization. An organization may choose to block file transfers, Webcam functions, or screen-sharing functions for IM communications. These types of actions will not prevent IM security breaches, but they could limit data loss. As with any policy and risk management, proper audit, reporting, and compliance controls must be in place.

35.4.2.6 Providing Secure Instant Messaging. In environments where IM is needed to run the business, the best strategy is to provide secure, managed IM services to the employees. Of course, the needs will vary among different organizations for different levels of IM connectivity, functions, and software. Many of today's popular corporate email and collaboration systems have built-in or optional IM services. When properly deployed, these IM systems can meet many of these secure IM best practices:

- Encrypt IM communications wherever possible: client to server, server to Internet, and so on.

- Encrypt logs and chat conversations at the workstation and server.
- Ensure that all logs, chat conversations, file transfers, and archives meet data transmission, retention, and destruction policies.
- Ensure that "presence awareness" features (software features that allow the user to communicate his or her presence or availability, such as "online," "away," or "out to lunch" to all users) comply with corporate personnel policies.
- Administratively disable features that cannot be encrypted or properly managed (screen sharing, file transfer, whiteboard, etc.).
- Where possible, lock or force configuration settings to ensure policy compliance.
- Establish procedures for periodic monitoring and auditing of IM systems; do not ignore logs.
- Enforce prudent password policies for IM systems.
- Properly secure IM communication systems with Internet connectivity; consider using proxies or intermediary gateways that protect internal corporate IM systems; ensure appropriate server lock-down policies and procedures.

Corporate-owned and -managed IM systems may not be possible in all situations. In those cases, the organization must form policies and procedures to limit risk and exposure with commercial IM systems. Some systems do provide "secure" IM, but one must be skeptical of exactly how much protection they provide. Consider limiting commercial IM needs to nonessential computers with limited network access, limiting or restricting users to specific IM applications or services, and monitoring instant message network traffic and usage. Some commercial IM services also provide "corporate" or "business" IM services, often for a fee. These premium offerings may provide the organization with the necessary or acceptable level of functionality and security. Access to public IM services should also *always* be blocked from machines with privileged access to critical data assets.

35.4.3 Response. Instant messaging breaches and compromised systems generally do not require special handling after a security incident. In general, normal policies and procedures can be followed to properly investigate, document, and respond to security breaches. There are many commercial tools, including forensic software, to aid in incident response. Infected or compromised systems, if no longer needed for investigation should be reimaged before redeployment to an employee; never allow a machine to be "cleaned" from an incident and returned to production.

35.4.4 Safe Messaging. Although most users at the organization are generally satisfied with mainstream IM systems, clients, and services, there are dangers to be considered. There seems to be almost an unlimited number of open source IM clients, Web-based IM and chat providers, social networking Websites providing IM, and the like. When considering policy and management of IM within the organization, it is important to analyze the source and intentions of all of the possible services. All IM software and services are not created equal; some may originate from untrusted sources and may contain malware and other security risks such as password stealers or keyloggers. Instant messaging software or providers may also be remotely logging information without the user's knowledge or consent.

Also, if the organization will utilize commercial IM software and services, it is critical to carefully examine the provider's terms of use and license agreements. The

responsibilities and liabilities of both parties should be carefully weighed by information security managers, company executives, and legal counsel before allowing use of the software and associated services.

35.5 SECURING SMS. Few technologies are more ubiquitous than SMS. Virtually all mobile phones are capable of sending and receiving SMS communications. Since mobile phones are virtually everywhere, security considerations must be in place to guard against the threats that they present. A technology with a relatively minor footprint can cause a world of destruction when used as a weapon. Today, SMS technology, and its associated complementary services, has grown exponentially. Securing and defending against SMS must be included in every organization's comprehensive security plan.

To understand SMS security, it is also important to look beyond the cell phone. SMS does not require a cell phone to utilize the technology. Many phone carriers allow SMS messages to be generated and sent from an unsecured, public Website. SMS messages may also originate from email messages, instant messaging services, and the like. SMS messages can even come from subscription services such as a daily horoscope or critical systems like emergency notification services. Beyond this, today's mobile phones, including smart phones, are more powerful and contain many more features and show no sign of slowing down. Phones are increasingly gaining processing power, memory, complex operating systems, and other features that essentially could redefine the device as a personal computer. Phones are able to access the Internet, install applications, communicate from phone to phone, and even access corporate data networks. Information security managers and professionals should never underestimate the power or versatility of a mobile phone. They are a threat to all of an organization's information security.

35.5.1 Dangers to the Business. SMS can introduce many types of security threats into an organization. SMS can cause a data breach by innocent mistakes or by deliberate attacks. This technology can be used as a criminal tool to deliberately steal information, to extract data, to extort information, and to deceive. It may also be a conduit for inadvertent data loss. The consequences of data lost via SMS are relatively the same as any other data breach: loss of confidence in the organization, loss of image, bad public relations, financial penalties, and so on. A serious or even minor data breach may appear to communicate to the world that the organization does not have a comprehensive security plan in effect, or the company does not abide by such a plan—whether true or not. Some investors, customers, or people in the general community may look at a breach of such a simple technology and ask, "How could the company not have proper security for something as simple as a cell phone?" A missing laptop with confidential data is a serious security breach, but a mobile phone, with all its capabilities, must be treated as nearly the same type of critical infraction.

35.5.1.1 SMS as a Tool for Deliberate Data Loss. One danger an organization may face involves an individual or group of people utilizing SMS technology to ferry critical data to unauthorized persons, usually outside the organization. This action would replicate an age-old tactic of stealing information piece by piece from within the organization to someone who should not possess the information. Consider classic tricks of criminals, such as copying information in tiny pieces over great amounts of time to avoid causing suspicion. Any number of technologies can be used to move data, including flash or thumb drives, iPods, scraps of paper, photographs, screen printouts,

embedded code, or even memorization. Disgruntled employees may use SMS to send confidential information to an accomplice or even to themselves for later use, such as selling the data, extortion, and the like. It would be virtually impossible to know that an employee is slowly leaking data outside the business from a mobile phone, especially when that phone is not owned or controlled by the organization. What may appear to coworkers as a serious text-messaging addiction may actually be a serious data breach.

Another fact information security management must consider is that SMS service, whether exactly true to the original definition or not, has expanded well beyond messages of only 160 characters. Mobile phone users are able to send real-time video streams, recorded video, photographs, substantially longer text messages beyond 160 characters, Web page links, and just about anything else the phone carriers can implement. If the mobile phone industry considers all of these features to be synonymous with SMS, the organization's security plan should as well. Business risk has increased greatly with every new technology addition.

35.5.1.2 Inadvertent Data Loss via SMS. Data loss can occur by mistake, bad luck, stupidity, misinformed user, or misunderstanding of features as well as by theft of the data device itself. Both deliberate data theft and inadvertent data losses are extremely dangerous, with potentially serious consequences. Search engines reveal many different tactics and war stories of data loss from a mobile phone as well as other SMS-specific security issues. These are scenarios and techniques to consider:

- SMS via email or the Internet
- SMS snooping or sniffing
- Recovery of improperly deleted data
- Stolen, mixed up, or lost phones
- Misdialed numbers
- Wi-Fi connectivity
- Unattended phone with no password
- Malware installed on phone (keyloggers)
- Recipient's phone is lost, stolen, or borrowed
- Impersonation

35.5.2 Prevention and Mitigation. SMS technology is not going to disappear anytime soon, so every organization must come up with a plan to prevent data loss and protect itself from this risk. Once again, the organization's leadership, security management, and security professionals must evaluate the risk of SMS versus the need to operate the business and maintain an amiable group of employees. Every organization must decide for itself exactly what kind of practices to put into place for SMS security and the cost/benefit of each practice. Everything must be considered, from policy and procedures, to deployment of security technologies and mobile phone company-provided services. With today's varying needs, newly emerging technologies and an array of mobile phones, it is difficult for any two organizations to adopt the same prevention and mitigation strategies. However, the next suggestions can be used to begin, update, or enhance the organization's security plan when addressing SMS technology.

35.5.2.1 Policy. An SMS policy should be developed for all cell phones that are company owned or company sponsored. The policy, when written on a solid foundation, is key for enforcement by human resources, especially when something may end up in litigation. Vague acceptable use policy will not be enough. A clearly stated position must be written, adopted, and communicated to all employees. The policy should apply to every employee, executive or trainee, with no exceptions. The policy should regularly be reviewed, updated, and redistributed, with recurrent training as necessary, especially in a rapidly changing world such as mobile technology.

A good policy must also address an important distinction common to mobile phone use in the organization: personal phones versus company-provided/sponsored phones. The policy must address: what is acceptable for employee conduct on the job; whether personal phones are allowed on the premises; what type of phone is allowed (usually refers to whether employees are allowed to have smartphones); where, when, and for what purposes can they use personal mobile phones; what is allowed on business-provided phones; and the like.

35.5.2.2 Mobile Phone Ban. In some cases, security needs may necessitate prohibiting the use or even possession of mobile phones on company grounds or in certain areas. This type of action should be included in company policy and should be clearly communicated. It may be necessary to remind employees with signs and repeated communication as well. This is a common practice to prevent data loss from any mobile phone function, including SMS. The organization should be sure to make distinctions for employee-owned phones and emergencies. If an area requires a high amount of security, err on the side of caution, and forbid mobile phones completely. The policy must extend to visitors, vendors, contractors, and other outside entities as well as to every employee—regardless of rank.

35.5.2.3 Providing Secure SMS. Providing "secure" SMS can prove to be difficult, and it can be easy to fall into a false sense of security. Information security managers must know exactly how their mobile phone infrastructure works before declaring the system secure. Although one component of a phone's connection may be secure—for example, from the phone to its messaging infrastructure—the entire path of an SMS message may not be secure. Some devices, such as the BlackBerry from Research in Motion, provide encrypted transport for messaging. However, email or messages to users on different phone networks or other messaging servers may not be encrypted. Information security professionals must clearly understand the technology they are deploying, and they must test for proper installation and configuration as well as working with the vendor to verify marketing claims. However, it is important to remember that if a solution is not as secure as the organization's policies and needs require, SMS, mobile phones, or both should be banned. Some phones or smart phone solutions allow administrators to "block out" services such as SMS or to install secure communications software. Carefully consider and evaluate all options and solutions.

35.5.2.4 Bring Your Own Device. There was a time when the corporate standard was the BlackBerry, provided by Research in Motion. While still popular and the most manageable platform for smartphones, a much more recent problem has emerged in today's organization. Tax law changes have increasingly made company-provided cell phones a thing of the past. Today, the rule is bring-your-own device; more

and more often, that device is an unmanaged Apple iPhone or Google Android variant smartphone.

Employees will want to connect their smartphone to the organization's email and calendaring system, as well as utilize the Wi-Fi infrastructure. They will want to stay connected for professional reasons (corporate email keeps them connected to the office and more productive), connected for personal reasons (personal email, social networking), and utilize battery-saving Wi-Fi instead of the cellular data networks. On the flipside, they may want to use their phone to tether another wireless device via Wi-Fi where they cannot or are not allowed to connect to corporate Wi-Fi.

This, combined with the knowledge these devices are unmanaged, have almost unlimited unknown applications, or "apps," and basically are small mobile computers, is enough to make most security pros wince. It is increasingly difficult to tell employees "no" when it comes to them utilizing these devices in the workplace, especially when they are sometimes required to do so for their job. This new challenge must be addressed in the organization's policy for mobile phones/mobile computing. It should be strictly laid out what is allowed, what is not, and how the rules will be enforced. Device management features should be utilized when allowed by the organization's platform to control things like encryption, email sync, data retention, device passwords, and device lock-out.

35.5.2.5 Other Considerations. The next list provides points to consider when planning for SMS security, many of which are from NIST Special Publication 800-48, "Wireless Network Security."[11]

1. Create policies and procedures to deal with lost mobile phones. The phone may contain sensitive data, including stored and deleted SMS messages.
2. If cell phones are banned from the organization's premises, ensure that physical security has procedures and rules for checking visitors and employees for mobile phones.
3. Many mobile phones have the capability to back up and synchronize their contents to the desktop. Ensure proper procedures to secure data and data leakage.
4. Policies and procedures should be in place to limit and manage the acquisition of mobile phones by employees—information security may not be aware of the existence of new phones in the environment.
5. Mobile phones are not easily audited, nor is there much software to aid in the auditing process.
6. Despite proper labeling of a company-owned device, if lost it will rarely be returned to the organization. Plan to mitigate damage caused by a lost mobile phone; utilize security features such as remote wiping the device after loss via "poison pill" features or "auto-destruct" features after several invalid password attempts.
7. If a mobile device supports screen-lock and power-on passwords, use these simple protections wherever possible.
8. Through policy and education, prevent as much sensitive and private information on the organization's mobile phones as possible, including SMS messages.
9. Utilize Public Key Infrastructure (PKI) technology where possible.
10. Install antivirus software where possible.

11. Utilize VPN and firewall technology for safer data communications.
12. If a phone is to be carried on international travel, SMS messaging should be prohibited if it all possible. The risks associated with taking a mobile phone to international destinations increase exponentially.

35.5.3 Reaction and Response. When a security incident involving a mobile phone and SMS does happen, unless the organization has in-house staff trained in mobile incident response or forensics, it may be best to work with the mobile phone provider, possibly also the manufacturer. Procedures and correct processes associated with data retrieval, preservation, investigation, and so on are best handled by those most qualified. If necessary, involve law enforcement. This is an area where a long-standing good relationship with local, state, or federal law enforcement is extremely beneficial—even if the investigation would not necessarily require law enforcement investigation. For more information on this subject, see Chapter 61 of this *Handbook*.

Investigating SMS issues, including tracking messages, tracking phone location, and tracking the path an SMS message took, can often be accomplished with the help of the mobile phone carrier. Law enforcement and court-ordered subpoenas may be necessary, depending on the situation.

Compromised devices should be carefully reviewed before returning them to regular use. Mobile phone providers can assist in "wiping" the device clean of all software, including malware if the device doesn't have such capability built-in. Specific practices and procedures vary by phone and provider, but some organizations may also choose to archive or destroy devices involved with a security breach of any kind.

35.6 SECURING COLLABORATION TOOLS. Information systems that provide online facilities for collaboration are increasingly valuable business tools. Although these tools provide excellent conduits for increased information sharing, they also have the potential to increase security threats. Even the Internet itself, with many Websites dedicated to information sharing, groupware, shared tools, and data storage, has become a collaboration workspace. New features and movements such as Google Apps, "Web 2.0," "The Cloud," and even free or online conference calling services must be taken into account in any organization's security plan. The nature of collaboration and the need to get critical business done efficiently is critical to most of today's organizations. Many companies and organizations are trying to get more work done with less people. Technology has become an important partner to allow employees to work together and to accomplish more in less time. Collaboration tools have become even more critical as businesses expand to include people working together from different geographical locations.

35.6.1 Security versus Openness. One of the longtime battles for security managers is security versus openness or functionality. The nature of collaboration requires uninhibited data and information sharing, which can be difficult to secure. Organizations have to find the right balance between allowing users free and open information exchange and providing the required level of security. Finding this balance takes cooperation and respect between the two groups: those who use the tools and those charged with securing the organization. The two groups must fully understand each other's position; without this understanding, finding a middle ground and negotiating compromise cannot take place. The goal of the organization surely must be efficient, uninterrupted business, but not at the expense of good security. The only way to work

through this complication is with good-natured, open, goal-oriented communication. This is *not* an information technology–only problem or process. Finding that optimum balance of security and functionality will require all types of management and staff to work together. Although this may be true of all information security domains, it is especially true of collaboration tools security. Without this important balance, the tools are essentially worthless: too secure and they will not be used, too open and the business could suffer catastrophic data and integrity loss. Some businesses are not able to recover from such a loss.

35.6.2 Dangers of Collaboration Tools. Collaboration tools are becoming powerful, and they must be given full security considerations. These tools should not be installed, accessed online, or integrated into the business without the proper planning, risk analysis, security configurations, and testing; ad hoc, unmanaged systems, installed without the knowledge of security personnel, must be prohibited, and violations corrected. Collaboration tools can easily become a nightmare for security management, especially if securing these tools is not a primary consideration from the beginning. Designing and implementing security measures on an already-deployed production system is invariably a frustrating exercise in futility for both the users and the information security personnel. Likewise, finding out after-the-fact that a cloud-based collaboration tool has become mission critical to the organization is a terrifying thought for most security professionals.

Some of the features and general dangers associated with many of these systems include loss of confidentiality, integrity, or availability. These dangers can occur due to any of these problems:

- **Lack of authentication requirements, rules, or procedures.** A wide-open system or one with poor authentication would allow for unauthorized persons to gain access.
- **Data snooping or capture.** Transmission of data to and from the system could be intercepted by unknown persons.
- **Impersonation.** Proving exactly who the user is may be difficult if not well managed, especially with weak authentication and authorization methods.
- **Unauthorized posting** of confidential information into insecure or public areas.
- **Misconfiguration.** A simple mistake in configuration could reveal private information.
- **Search engines.** Documents or other information may be subject to search engine crawlers/agents/spiders if proper security is not established.
- **Rogue collaboration systems.** If a department or group deploys its own tools, privately or publicly, without the knowledge of the security group, proper security cannot be guaranteed.
- **Internal threats.** One cannot be concerned only with external threats. One department's collaboration system may be another department's limitless temptation.
- **Users.** Users may not always have security in mind. Small mistakes or shortcuts could lead to major security breaches.

When deploying or evaluating collaboration tools, risk analysis must be performed to determine if the organization is able and willing to accept the associated risks. Security groups should thoroughly brainstorm and research as many possible security

threats to the collaboration system as possible. It may be beneficial to work with the solution provider's support group to minimize or eliminate as many security risks as possible.

Workgroups utilizing collaboration tools place a great amount of trust in the application and the tools. Many of the applications available today are light on security and heavy on marketable features. Although many online companies have become much more serious about data security, they are not the owners or protectors of the organization's data; that is still up to the organization.

35.6.3 Prevention and Mitigation. Collaboration tools and systems should receive the same security care as any other information system. Although the nature of collaboration may be somewhat open, the same policies, procedures, and careful controls should apply. The goal must still be the confidentiality, integrity, and availability of the data and the information system. Collaboration tools must still benefit the business while ensuring the business will not be harmed by a security incident. By taking the necessary steps to prevent and mitigate security issues, collaboration tools can be invaluable to the organization.

The next suggestions can be utilized to aid an organization in securing collaboration tools and systems.

35.6.3.1 Policy. It is debatable whether collaboration tools necessitate specific, separate policies. What is more important is that a complete, well-written, and well-communicated policy exists, one that includes provisions for collaboration tools, systems, and associated technology. Clear understanding and communication of collaboration tool security must be well researched, well written, concise, well communicated, and updated regularly. It is critical that the policy remain valid as new and more complex collaboration tools are developed and deployed.

Policies should also include security options that may otherwise be out of the control of the organization. For example, if a company forbids using public file-sharing services, the policy should cover users attempting to use the service from outside the organization as well as within it. Employees should not be able to use services or systems that do not comply with the policy, no matter where or how the service is to be used.

A good policy should be inclusive, especially when defining exactly what the organization considers a collaboration tool. It would be easy to forget applications such as email, IM, online meetings, blogs, social networking, shared network resources, remote access software, peer-to-peer file sharing, and the like. Many technologies have collaboration components that must be considered to ensure security.

35.6.3.2 Prevent Access or Use. Another option, in conjunction with policy, is to block the use of collaboration tools, depending on the organization's needs. This may involve deploying technology to accomplish this goal, including content blocking, firewalls, or both. This should disallow installation or use of rogue collaboration tools. Periodic review of networked systems and network traffic should be conducted to ensure compliance with prevention or limitation of collaboration tools.

35.6.3.3 Limit Access. Many collaboration tools can be deployed as an internal-only system, a cloud-based system, or both. Organizations will want to choose how users will access these systems. For example, disallowing unsecured communications from the Internet may help increase security. Likewise, it may be necessary to block

access to public services from within the organization's network. Or technical solutions, such as VPN connections from outside the organization's network, may be used to meet communication needs.

35.6.3.4 Deploy or Enhance Security Frameworks and Technologies. Wherever possible, install solutions that will increase collaboration tool security and that can be integrated into existing security frameworks. If the organization has a high-security, single sign-on solution, integrate the collaboration systems into it. Another example would be to integrate the collaboration systems into a new or existing PKI infrastructure. Utilize well-known and reliable solutions such as Secure Sockets Layer (SSL) and encryption for the host and all participants. This greatly reduces the risk of security breaches during data transmission.

35.6.3.5 Audit. No matter what level of policy, procedures, or preventions are put into place, every organization *must* audit for compliance. Procedures for auditing collaboration tools and their use should be included in the organization's regular, structured, information security auditing functions. Any deviations from the policies and procedures mandated for collaboration tools must be acted on in a timely manner.

35.6.3.6 Monitoring. Any organization that deploys collaboration tools must monitor and report on the system's usage, audit results, and data contents. (The organization must examine the actual data contents to ensure compliance with protections such as protected health information [PHI] or Social Security number [SSN]. Many new products have rules written for this reason.) Monitoring and reporting work to ensure that collaboration tools and systems are being used for their intended purposes. Monitoring and reporting of active projects should look for unusual patterns of use, policy violations, inactive users, inactive or outdated systems, and the like. Proper system management should already be in place, but it is important to check the systems periodically. For example, if a group is utilizing a collaboration system for a project, once the project has been completed, all project materials and users should be removed from the system. Reports from system monitoring and auditing should be acted on at once.

35.6.3.7 Consider Outsourcing and the Cloud Carefully. Some organizations are tempted to use commercial online-only collaboration systems or hosted solutions. This decision should not be made lightly; consider the risks versus the returns. The organization should carefully review all terms of service, license agreements, service-level agreements, and legal responsibilities carefully. Legal counsel must be involved to ensure the organization is protected, especially in the area of data ownership, possession, legal discovery, and subpoena power.[12]

The cloud has gotten a *lot* of attention lately. While cloud security is far beyond the scope of this chapter, the cloud should be treated no differently than any other outsourced application. The cloud often contains the same risks as any other service where the data is stored outside of the control of the organization. Security should be a primary concern when evaluating solutions in the cloud, with consideration given to:

- Service level agreements
- Data ownership

- Data security technology in use while data is at rest (encryption, backup, replication, etc)
- Data recovery if the provider is sold, closes, or goes bankrupt
- Physical location of the data and what national, state or local laws apply to that location and data
- Administrator access (for the customer and which administrators at the service provider also have access)
- Liability of data breeches at the service provider
- Key escrow for encrypted data

While this is not an exhaustive list, this should help an organization in brainstorming and researching the risks of cloud-based collaboration tools. There are a lot of online resources from reputable organizations that have already provided guidance for organizations wading into the cloud. As usual, consult legal counsel.

35.6.3.8 Audits and Penetration Testing. As with most information systems, providing necessary security should involve regular, external, third-party penetration testing and audits. Collaboration tools and associated systems should be tested and evaluated for their security fitness. Any problems discovered should be documented and swiftly remedied. Allowing a neutral, external entity to test the system independently is superior to internal testing, so that bias can be ruled out.

35.6.3.9 Keep Collaboration Tools Current. Keeping collaboration tools and their associated information systems up to date is critically important. Applying patches for vulnerabilities is good information technology and information security best practice. After thoroughly testing patches in a test environment, they should be applied to production environments as soon as possible. Do not ignore software vendor patches, especially those for known vulnerabilities.

35.6.4 Reaction and Response. Once a security breach has been discovered involving collaboration tools, the organization's usual policies and procedures should be followed. Procedures for compromised information systems should be well-formed, repeatable processes to preserve evidence, provide for rapid discovery and investigation, and meet necessary regulatory guidelines. When necessary, law enforcement, legal advisors, or both should be utilized to ensure proper evidence collection and documentation. Policies should also dictate the procedures for post-investigation tasks as well, such as requiring compromised systems to be copied, archived, destroyed, reimaged, or reinstalled. It is generally not advisable to try simply to "clean up" a compromised system. It can be difficult to guarantee that a compromised system once again has integrity.

35.7 CONCLUSIONS. This chapter introduces security managers and professionals to securing peer-to-peer technologies, instant messaging, short messaging services, and collaboration tools. The suggestions and information in this chapter are meant to aid in making decisions regarding these tools within the organization's overall security plan. Many of the examples and concepts are meant to aid in the planning, policy development, and review of the organization's exposure to these technologies and their dangers. This chapter should serve as only a starting point for the

organization's research on each topic and to ensure that information security managers at least have a brief understanding of each concept, its risks, prevention and mitigation strategies, and suggestions for response. It is difficult to recommend solutions for every type of business, so each organization must make its own judgment for securing these technologies. The popularity and ubiquity of P2P, IM, SMS, and collaboration tools ensures that they will be part of every security plan for many years to come.

35.8 FURTHER READING

Covill, Jared J. *Going Google: Powerful Tools for 21st Century Learning*. Corwin, 2012.

Flynn, Nancy. *Instant Messaging Rules: A Business Guide to Managing Policies, Security, and Legal Issues for Safe IM Communication*. AMACON, 2004.

Kunz, T., and S. S. Ravi, eds. "*Ad-Hoc, Mobile, and Wireless Networks*." 5th International Conference, ADHOC-NOW 2006, Ottawa, Canada, August 17–19; 2006 Proceedings. New York: Springer, 2007.

Lamont, Ian. *Google Drive & Docs in 30 Minutes*. Digital Media Machine, 2012.

Piccard, P., B. Baskin, G. Spillman, and M. Sachs. *Securing IM and P2P Applications for the Enterprise*. Norwell, MA: Syngress, 2005.

Rittinghouse, J., and J. F. Ransome. *IM Instant Messaging Security*. Burlington, MA: Elsevier/Digital Press, 2005.

Rittinghouse, J., and James Ransome. *Cloud Computing: Implementation, Management, and Security*. CRC Press, 2009.

Rhoton, John, Jan De Clercq, and David Graves. *Cloud Computing Protected: Security Assessment Handbook*. Recursive, 2013.

Taylor, I. J., and A. Harrison. *From P2P to Web Services and Grids: Peers in a Client/Server World*. New York: Springer, 2004.

Winkler, J. R. *Securing the Cloud: Cloud Computer Security Techniques and Tactics*. Norwell, MA: Syngress, 2011.

35.9 NOTES

1. Dropbox, "How Secure is Dropbox?" Dropbox Website. 2013, https://www.dropbox.com/help/27/en
2. Michael Kassner, "DropSmack: Using Dropbox to Steal Files and Deliver Malware," *TechRepublic | Security*, April 15, 2013, www.techrepublic.com/blog/security/dropsmack-using-dropbox-to-steal-files-and-deliver-malware/9332
3. M. E. Kabay, "Dropping the Ball on Dropbox," *InfoSec Perception*, April 19, 2013, http://resources.infosecskills.com/perception/dropping-the-ball-on-dropbox/
4. Matthew J. Schwartz, "5 Dropbox Security Warnings for Businesses." *InformationWeek/Security*, August 14, 2012, www.informationweek.com/security/management/5-dropbox-security-warnings-for-business/240005413?pgno=1
5. Cloudfogger, "Protect Your Privacy on SkyDrive, Dropbox, Google Drive in the Cloud," *cloudfogger*, 2012, www.cloudfogger.com/en
6. It should be noted, that at the time this chapter was updated, it is not uncommon for a house or small organization to have an Internet connection of 40–50 mb/sec. While the resource-hogging nature of P2P is becoming less and less of an issue, that should not negate the need to have a policy to limit this technology in the organization.

7. J. Borland, "'Spyware' Piggybacks on Napster Rivals," *CNET News.com*, May 14, 2001, http://news.cnet.com/2100-1023-257592.html
8. N. Hindocha, "Instant Insecurity: Security Issues of Instant Messaging," *SecurityFocus*, January 13, 2003, www.securityfocus.com/infocus/1657
9. P. Festa, "ICQ Logs Spark Corporate Nightmare," *CNET News.com*, March 15, 2001, http://news.cnet.com/2100-1023-254173.html?legacy=cnet
10. National Vulnerability Database, http://nvd.nist.gov/nvd.cfm
11. K. Scarfone and D. Dicoi, *Wireless Network Security for IEEE 802.11 a/b/g and Bluetooth (DRAFT)*, NIST Special Publication 800-48 Revision 1 (Draft), 2007, http://csrc.nist.gov/publications/drafts/800-48-rev1/Draft-SP800-48r1.pdf (URL inactive).
12. M. Rasch, "Don't Be Evil," *SecurityFocus*, 2007, www.securityfocus.com/print/columnists/447 (URL inactive).

CHAPTER 36

SECURING STORED DATA

David J. Johnson, Nicholas Takacs, Jennifer Hadley, and M. E. Kabay

36.1 INTRODUCTION TO SECURING STORED DATA 36·1

36.1.1 Security Basics for Storage Administrators 36·2
36.1.2 Best Practices 36·2
36.1.3 DAS, NAS, and SAN 36·4
36.1.4 Out-of-Band and In-Band Storage Management 36·4
36.1.5 File System Access Controls 36·5
36.1.6 Backup and Restore System Controls 36·5
36.1.7 Protecting Management Interfaces 36·6

36.2 FIBRE CHANNEL WEAKNESS AND EXPLOITS 36·7

36.2.1 Man-in-the-Middle Attacks 36·7
36.2.2 Session Hijacking 36·8
36.2.3 Name-Server Corruption 36·8
36.2.4 Fibre-Channel Security 36·8

36.3 NFS WEAKNESSES AND EXPLOITS 36·8

36.3.1 User and File Permissions 36·9
36.3.2 Trusted Hosts 36·9
36.3.3 Buffer Overflows 36·9

36.4 CIFS EXPLOITS 36·9

36.4.1 Authentication 36·10
36.4.2 Rogue or Counterfeit Hosts 36·10

36.5 ENCRYPTION 36·10

36.5.1 Recoverability 36·11
36.5.2 File Encryption 36·11
36.5.3 Volume Encryption and Encrypting File Systems 36·11
36.5.4 Full Disk Encryption 36·12
36.5.5 Vulnerability of Volume, File System, and Full Disk Encryption 36·12
36.5.6 Database Encryption 36·13
36.5.7 Smart Phone Encryption 36·15

36.6 DATA DISPOSAL 36·15

36.7 CONCLUDING REMARKS 36·16

36.8 FURTHER READING 36·16

36.9 NOTES 36·17

36.1 INTRODUCTION TO SECURING STORED DATA. This chapter reviews methods of securing data stored on nonvolatile media. Nonvolatile media include magnetic disks and their (hard) drives, compact discs (CDs), and digital video disks (DVDs), with their optical drives), and flash drives (also known as *USB drives*, *flash*

disks, and *memory keys*). Volatile storage devices, which are not covered in this chapter, include random access memory (RAM) and other storage that loses its contents with a power loss.

One of the most important issues for security management is the frequent loss or theft of laptop computers and other mobile computing devices. For example,

- A 2008 report suggested that in the United States alone, over 600,000 laptop computers were left behind at U.S. airports every year, with about two-thirds of them left unclaimed.[1]
- In 2010, "… 275 businesses in Europe… lost a combined 72,000 laptops, costing those organizations a total of $1.8 billion. Of the 265 laptops lost by the average organization per year, a company will typically only recover 12."[2]
- In 2012, the rate of loss was still 12,000 laptops lost per week at U.S. airports; another detail of reports at that time was that the unclaimed laptop computers were being auctioned to anyone able to pay for them.[3]

Unencrypted files on these stolen or repurchased laptop computers are an invitation to violations of confidentiality and control. For example, an audit of discarded computers from the New Jersey state government revealed that "Taxpayers' Social Security numbers, confidential child abuse reports, and personnel reviews of New Jersey workers nearly went to the highest bidder after the state sent surplus computers out for auction." In addition, "Nearly 80 percent of discarded computers in a comptroller's office sample had not been scrubbed of data before being shipped to a warehouse. .."[4]

Since the last edition of this *Handbook* was published in 2008, cloud-based storage solutions are increasingly being used by individuals and by organizations, so this chapter includes a short discussion of security for such repositories.

36.1.1 Security Basics for Storage Administrators. Storage systems have developed outside the security umbrella of other organizational assets. Because the storage arena is one of the most strategic parts of the infrastructure, professionals should take the same care in developing comprehensive security controls as those that are addressed for the remainder of the network. A number of vendors have focused on the development of secure storage environments that are scalable yet flexible; most address both the logical and physical aspects of security. However, any appropriate strategy for data storage protection includes a balance between protecting the confidentiality and integrity of the information while also ensuring its availability and utility to the system and to authorized users. Ultimately, those with a responsibility for data storage will also be tasked with maintaining this balance at a reasonable cost.

Knowing how and where data will be stored on a network and addressing the known risks to the data is the best option, as it is often more efficient in terms of resource usage (time and money). An organization need not be a high-profile entity to suffer from a compromised pool of data. A single backup can contain enough concentrated personal or sensitive corporate information to experience a loss of credibility, to lose revenue, and possibly to bring the organization to its knees. Worse, a backup that is copied illicitly may show no signs of having led to loss of control over confidential data.

36.1.2 Best Practices. Every organization must keep its applications, servers, and end-user systems up and running to make use of information and to maintain the highest degree of information availability and integrity. A tiered data protection model

works best; it includes a layered defense, due diligence, and restricted management. Implemented correctly, security should be transparent. Best practices for providing a secure data storage environment include the following:

- Performing an audit and risk assessment on the storage infrastructure, looking for risks and vulnerabilities.
- Implementing authentication across the storage network that could coordinate authorization, password maintenance, and encryption.
- Implementing strong role-based access controls and assigning access rights to parties on a need-to-know basis.
- Adopting and enforcing data-encryption and data-classification policies. Based on the classification level assigned to the data, the organization's policy may require the encryption of the data at rest throughout the lifecycle of the data. There may also be requirements to encrypt the data in flight (across the network) as well.
- Requiring strong security features and practices from storage-system vendors and offsite-storage providers.
- Remembering to secure the Storage Area Network (SAN) at the switch (or fabric) level. Carving up the fabric by zones is one technique that limits access to various parts of the SAN.
- Including any data-replication or storage-replication technologies in the overall storage security plan, where such replication traffic may include transaction journals and other temporary locations with partial or full copies of sensitive data.
- Evaluating and implementing data-loss prevention (DLP) technologies to maintain the security of stored data outside the enterprise. Data tagging, or labeling, along with pattern matching (e.g. 999-99-9999 for a Social Security number) are key for assisting with DLP.
- Creating a policy for discarding old devices and media and data-storage services, to include routinely performing tasks such as secure wiping or physically destroying all data-storage devices and media.
- Evaluating and implementing retention and data-destruction policies, ensuring compliance with any applicable organizational or government regulatory issues.
- Isolating the storage management network from the organization's primary network. By not isolating the network, every employee potentially has access to the stored data.
- Establishing access-log monitoring.
- Performing employee and contractor background checks as part of the Human Resources (HR) hiring procedures.
- Establishing facilities controls in the organization to restrict physical access to data centers, locking storage cabinets and server racks, using locks built into some servers, and ensuring the reliability of the perimeter and building(s) and verifying that such controls are in place for remote storage locations.
- Treating the security of backups as warranting monitoring and alarms at a high level of importance. Adopting secure media management tracking and handling policies that include backup requirements for financial information, employee data, and intellectual property.[5] Chapter 57 in this *Handbook* contains much information about data backup.

36.1.3 DAS, NAS, and SAN. There are three primary methods for storing data: Direct Attached Storage (DAS), Network Attached Storage (NAS), and Storage Area Networks (SANs). In addition, remote data storage using cloud computing services are growing in popularity, especially with individual users but even with organizations.

Direct Attached Storage drives are those that are connected directly to the computer. DAS can either be internal, contained within the computer's case, or external and attached via a peripheral component interconnect, PCI, eSATA, or other bus channel. The risks to DAS devices are either their physical theft or access through the computer system that they service.

Network Attached Storage devices are specialized servers that run minimized operating systems and file systems designed specifically to support input/output (I/O) from other servers. The servers that attach to the NAS devices have DAS that contains their operating systems, applications, and other components but, normally, write all data to the NAS device via TCP/IP over Ethernet connections. NAS is utilized via file sharing protocol such as Network File System (NFS) for UNIX systems and Server Message Block (SMB) or Common Internet File System (CIFS) for Microsoft systems. (As CIFS grew out of SMB, the two are often noted as SMB/CIFS or CIFS/SMB.) As with any Ethernet connection, connections between a system and the NAS server that it uses are subject to being sniffed, to eavesdropping, and to packet capture. NFS and CIFS threats are discussed later in this chapter.

Storage Area Networks are collections of centralized disks that can be accessed by numerous servers. Using SANs can facilitate company growth, as most SANs options allow additional disks to be added to the pool as data storage needs increase, without having to take the attached systems offline, as would need to occur if new DAS were being added to individual systems. Data backups can also be easier to control, as a single storage resource could potentially be backed-up instead of each individual system. With the implementation of RAID or other disk redundancy implementations, writing data to multiple disks can be accomplished without impacting the application servers' performance. Systems can be attached to the SAN by various methods including TCP/IP and fibre channels. Fibre channels are discussed later in this chapter. Using IP for SANs connections enables servers to connect over the Internet—but this option must be used with caution, due to the security concerns of transferring data over the Internet.

Cloud backups use Internet connectivity to store backup data on remote systems out of the control of the data owners.

36.1.4 Out-of-Band and In-Band Storage Management. Managers may have to control storage locations remotely, that is, from a location other than a directly attached console. There are two approaches to such control communications, each with its own particular security issues.

In-band management uses the same network as the data transfers; *out-of-band management* uses a separate network.[6]

While in-band storage management uses the same channels as the data itself traverses for storage, out-of-band management uses alternative methods. For example, an out-of-band solution might have a storage administrator working from a desk and connecting to the storage system over the primary network used by all employees, while the data traverses a dedicated channel between the application server and the storage system.

With out-of-band management, consideration must be given to how to ensure that only authorized systems, such as the administrator's, are connecting to the storage system. This is especially necessary if the storage system does not require authenticated

connections being established before a command is accepted. Without authentication, any system able to communicate with the storage system could issue commands that would negatively impact the storage system. Another risk is due to the interface used by management communications and the commands being sent across the wire without being encrypted. For storage systems that are managed by HTTP interfaces by default, it may be possible to use HTTPS instead, in order to mitigate the risk of commands and logins from network packet capture.

In-band management also has concerns. Commands sent in-band are normally sent in clear text. Other threats of in-band storage management include[7]:

- Management interfaces being subjected to denial-of-service (DoS) attacks,
- Commands providing information on other devices and controllers, and
- Set and reset commands being issued inappropriately.

36.1.5 File System Access Controls. File systems provide access control to data. UNIX file systems provide controls based on the user owner, the group owner, and "other," or those that are not the user or a member of the group that owns the data. Microsoft® Windows systems allow for data owners to be specified and access granted by either individual usernames or group accounts. Access control lists (ACLs) can also be used to provide access exceptions to the normal access permissions of the data files.

When correctly applied, these access controls can be effective in preventing unauthorized access to data through normal usage. However, the file system trusts the computer's operating-system access controls to have correctly authenticated and authorized the user. If the access controls for the operating system are circumvented, then file system access controls lose their effectiveness.

For more details of operating system and local area network access controls see Chapters 24 and 25 in this *Handbook*.

36.1.6 Backup and Restore System Controls. Systems used for the backup and restoration of data need extra security consideration because the data contained on them are often critical to disaster recovery and business continuity. There are several threats to backup data that are not faced by attacks against other data stores. While most systems only have their data stored on disk (DAS, NAS, or SAN), backup systems often write data to tape cartridges or other removable media specifically intended for off-site storage. Another option is to backup data electronically to a remote system, either at another of the organization's data centers or with an off-site backup vendor. As with any transmission of data to remote facilities, the data must be protected during transit and at the destination facility.

Regardless of the media used, all data backups need to be stored at a secure, environmentally protected facility sufficiently distant from the originating location to mitigate the risk of losing both primary and backup data from a single major event, such as an earthquake, flood, or volcanic eruption. Additionally, the backup media storage needs to be secured by restricting access to authorized personnel only. Media used for backups needs to be evaluated to ensure that it meets or exceeds the longevity needs of the data, as defined by the organization's backup policies and standards.

Additional risk from system imitation also needs to be considered. If an attacker is able to insert a system that impersonates a backup system, all of the data intended to be backed-up may instead be written directly to the attacker's system. Conversely, if an attacker is able to insert a system that can masquerade as one or more data storage

systems, then the attacker could request a restoration of data from existing backups, gaining unauthorized access to the information.

To mitigate such risks, interactions between data storage systems and backup systems should be authenticated. For manually controlled backups, systems could have backup accounts created that require an interactive login to authenticate the backup request. For automated backups, other options may be available, including the use of client and server certificates to authenticate both systems involved in the backup connection.

In recent years, developments in cloud computing technology and cloud-based storage have enabled organizations to leverage hybrid data-backup solutions. Instead of backing up data at regularly scheduled intervals, changes to systems replicate to a cloud-based service in near real-time, effectively acting as a replicated copy. The addition of de-duplication technology, sending only those blocks of data with changes, further improves the efficiency of the solution. However, as with all remote backup solutions, continual due diligence with the service provider is necessary to ensure that appropriate security controls exist for the transmission, storage, access, and retrieval of the data. This is especially important given the "on demand" nature of a cloud or hybrid-based solution for data or system recovery. Any use of cloud services should be done only after:

- A thorough assessment of the cloud service provider's (CSP) security practices,
- Validation that storage of the data in the solution is in compliance with all legal, regulatory, and company requirements, and
- Employing data encryption that maintains in-house management of the data encryption keys.

Off-site storage of data, whether written directly to a storage vendor's server or on removable media such as tape, deserves its own security considerations. For example, how does the vendor secure physical access to its site? How do they vet their employees? Most risks such as these can be mitigated by performing due diligence of the vendor prior to entering a contract, and by eliminating prospective vendors that do not meet the organization's security needs.

Another mitigation strategy is to encrypt data backups before they are sent offsite. Many backup applications offer encryption methods. The use of encryption for files and other data storage systems is discussed later in this chapter. Similar techniques can be applied to backup data by encrypting it as it is written to the backup media.

For more details of backup strategies and security, including cloud backups, see Chapter 57 in this *Handbook*.

36.1.7 Protecting Management Interfaces. Management interfaces pose one of the greatest threats to the security of stored data. These interfaces provide administrative access to the data stores, allowing individuals with appropriate access the ability to manipulate data elements, update account security, and perform other housekeeping activities. Therefore, care is required when implementing a storage solution to ensure that a well-defined defense in depth exists. While two-factor authentication is more secure, it is not always practical. At a bare minimum, each administrative user should have a set of credentials, and complex password requirements, with a regular password change frequency. For situations where the data storage occurs via an Internet-based source, certificate or token-based authentication can provide an extra

layer of authentication security. In addition, the individuals responsible for administering security should not be the same individuals responsible for managing the storage environment. This separation of duties limits the ability of any one individual to circumvent security controls, without some type of conspiracy between the storage and security administrators.

Another important component of the defense in depth strategy involves the use of audit logs. Logging should be enabled to detect violations of policy and of prescribed procedures. However, logs are valueless unless subjected to regular and random review, with follow-up if anomalies are detected. It is unrealistic to expect an individual to pore over voluminous log files on a daily basis. However, log aggregation and correlation technology can be employed to provide an additional layer of confidence. Regardless of the final implementation, the use of audit logs, and restrictions on the ability to access and modify those logs, plays an important part in guaranteeing that no data corruption occurred.

Mitigating these risks to the management interface requires careful monitoring and control of who can access and install these interfaces. With the movement to Web-based interfaces, this discussion comes down to strict user access control and authentication. In any case, it is imperative that only trusted users with a need-to-know be allowed access to the management interface. Once inside, individuals can manipulate the environment as needed to support their goals, whether to further organizational objectives or to perform malicious actions.

36.2 FIBRE CHANNEL WEAKNESS AND EXPLOITS.

Fibre channels, while very economical, present unique challenges to the storage environment. The term *Fibre channel* refers not just to an optical-fiber-based communication pathway, but rather to a complex communication protocol. There are a number of inherent weaknesses in the technology, some of which are fairly straightforward and manageable, but others introduce questions of the viability of the technology in larger storage environments.

One of the most serious security weaknesses with Fibre channel is that all communications occur in cleartext. Fortunately, when Fibre channel implementations occur completely within a data center or other secured area, this is not a major concern, as the ability of unauthorized individuals to intercept traffic *on the wire* is limited unless they have physical access to the cabling and can avoid detection as they attempt to intercept traffic or to install interception devices. In early Fibre-channel implementation, native authentication and encryption was not available. However, newer fabric switches and SAN devices provide built-in capabilities for authentication and encryption to significantly reduce the risk posed by cleartext data transmission.

From a vulnerability perspective, attackers can use the Internet Protocol (IP) to craft exploits against Fibre channel, since both protocols use a frame-based communication scheme. Unfortunately, based on the cleartext issue discussed above, an attacker could sniff frames from the Fibre channel connection, and gain information needed to craft an attack. This section focuses on three types of common attacks: *man in the middle*, *session hijacking*, and *name-server corruption*. Most attacks on a storage network require physical access to that environment, or access to the appropriate sniffing hardware, increasing the difficulty of successful, undetected attacks.

36.2.1 Man-in-the-Middle Attacks.

Man-in-the-middle attacks (MIMAs) take advantage of weaknesses in the frame-based communications through Fibre channel. Like IP-based attacks, MIMAs involve an attacker intercepting communications, stealing or changing data, and passing that frame on to the intended destination. Fibre

channel includes a *Sequence ID* and *Sequence Count*, both intended to ensure consistent communication from sender to receiver. Much like IP, the Sequence Count is an easily predictable sequential number, allowing an attacker to anticipate the next sequence number and forward a packet ahead of the sending system. The introduced packet allows the attacker to intercept the stream without the authorization of either party. Mitigating this issue requires data-integrity checking to guarantee reception of the correct information and rejection of the fraudulent packets.

36.2.2 Session Hijacking. Session hijacking presents the same type of problem as MIMAs, and occurs in much the same way. However, this type of attack focuses on the lack of authentication in Fibre-channel environments. Instead manipulating data in each frame and passing it on, the hacker uses knowledge of the Sequence ID and Sequence Count to intercept and control the session, making the recipient believe that the attacker is really the original sender. The newly controlled session can then be used to extract whatever data or other information the attacker wishes. Mitigating this issue requires strong authentication to provide a guarantee that the original sender is still the same system throughout the length of the connection.

36.2.3 Name-Server Corruption. The last type of attack in this section involves address spoofing, similar to DNS spoofing in the IP world. Each Fibre-channel connection registers its name with World Wide Name (WWN) service, through two processes, a *Fabric Login* (FLOGI) and a *Port Login* (PLOGI). Typically, name-server corruption occurs during the PLOGI process by allowing an incorrect host to register itself with the Fibre channel switch (which contains the WWN service) using a spoofed address. The switch registers the host under that address as if it were valid because of the lack of any host authentication. When the real host tries to connect, the switch denies that connection because the incorrect host is already connected. This type of attack requires some timing on the attacker's part, but can be easy to accomplish because of the weaknesses previously discussed. Modern fabric switches provide additional authentication mechanisms to defend against these attacks through validation of the device with the switch at regular intervals.

36.2.4 Fibre-Channel Security. This short examination of Fibre channels has pointed out that many weaknesses exist. However, this does not mean that Fibre channel should be dismissed as a suitable technology. When implementing Fibre channel into an environment, care must be taken to address these vulnerabilities, taking into consideration location, distance, and availability of the implementation to individuals, systems, and other devices on the network. The physical placement of the devices and wiring must also be optimized to mitigate the risks associated with using Fibre channels. Vendors are also answering the call by offering technology to help secure Fibre channel through built-in authentication and encryption capabilities.

36.3 NFS WEAKNESSES AND EXPLOITS. Network File Systems (NFS) provide a service allowing a user on a client machine to access network-based resources as if they were local to that user. This service is built upon the Remote Procedure Call (RPC) service, and although very useful in a networked computing environment, NFS presents a number of security issues that must be addressed prior to implementation. NFS is typically geared toward high-bandwidth environments, such as an LAN, or networks sharing nonsensitive information. Since NFS does not provide encryption

between hosts, using this technology for other networks, especially those exposed to the Internet, introduces additional risk.

This section describes three of the most common weaknesses and exploits for NFS: user and file permissions, trusted hosts, and buffer overflows.

36.3.1 User and File Permissions. Aside from the lack of encryption, NFS allows user access based on the particular host connected to the NFS share. This means that any user connected to that host can access the network resources. Restricting users to read-only access eliminates the potential to use NFS as a collaborative technology, because users can no longer create or update information on the shares.

When mounting shares with read-write access under NFS, any user connected to the host can access another user's files, as the only protection the file has is its permissions. Administrators attempt to mitigate this risk by forcing all users to access the share under a group or common set of read-only credentials, but this approach eliminates some of the benefits that a network share provides. The read-only share then requires administrators to update or edit files, providing a library-style approach rather than a collaborative file management environment.

36.3.2 Trusted Hosts. Problems with NFS specifically concern the authentication of hosts to the NFS environment. Because NFS controls mount requests based on the host connecting, and not on the particular user, a rogue host could request an NFS mount and make changes to resources. An attacker could also compromise a DNS server used by the system exporting the NFS to point that system to an unauthorized machine. Because there are no login credentials shared prior to mounting NFS shares, if the hosts are not trusted, there are no additional checks to validate the integrity of the new host.

36.3.3 Buffer Overflows. In many implementations of NFS, data input checking does not occur before processing a request. This flaw presents an opportunity for a buffer-overflow attack. When a directory-removal request comes to an NFS server from a user with read-write privileges, the server does not check the length of the path name, and the user can include additional instructions beyond what the server should receive. Those instructions, presumably malicious, could then be executed by the server as an administrative account, such as *root* or *administrator*. As a result, unintended or unauthorized data manipulation can occur.

Recent implementations of NFS have included Kerberos authentication to help validate the users and what they are able to do. In addition, the same validation can be used to validate hosts before they connect to the NFS server. However, these improvements are only partial solutions. Criminal hackers continue to develop and apply new buffer overflows, and it would be unwise to assume that NFS, or any other network technology, can be completely secured.

For more detailed discussions of secure programming techniques and software quality assurance to preclude buffer overflows and identify and prevent other software vulnerabilities, see Chapter 38 and 39 in this *Handbook*.

36.4 CIFS EXPLOITS. The Common Internet File System (CIFS), an Internet-enabled Server Message Block (SMB) protocol, builds upon that protocol by including encryption and secure authentication to the existing resource sharing capabilities of SMB. Unfortunately, from a security perspective CIFS blends some of the new with some of the old, and the result includes a number of security issues.

36.4.1 Authentication. CIFS implementations provide either a straight password-based authentication scheme or a challenge-response scheme. Both of these approaches occur in cleartext, allowing anyone with wire access to intercept and capture authentication credentials to the network share. Even with the challenge-response approach, attackers could spoof a transaction and gain access to the share. Recent implementations of CIFS rely on Kerberos for authentication, and much like NFS, while providing additional security, introduces Kerberos-based vulnerabilities, which are outside the scope of this section. Some CIFS implementations also provide a "share-level" security model, rather than a "user-level" security model. In essence, rather than each user maintaining individual credentials, the share has only one set of credentials, which all users share. The weaknesses inherent to a share-level model are similar to those found with the use of group or shared accounts on any other system.

In addition to the authentication issues, CIFS is also vulnerable to dictionary and brute-force attacks against a user's credentials. These generally involve a chosen plaintext attack, helped by intercepting challenge-response pairs during the authentication process. However, both online and offline dictionary attacks are available to the attacker based upon the amount of time available to watch the connection attempts.

36.4.2 Rogue or Counterfeit Hosts. It is important to identify the differences between the attack surface of CIFS and that of NFS. Man-in-the-middle and the trusted host issue both apply to CIFS, with some different concerns. Improperly configured CIFS clients can be fooled into thinking that they should supply a password instead of interacting with a challenge-response scenario, thus supporting man-in-the-middle attacks. Additionally, if a CIFS environment that does not enable session or message authentication, it removes security controls designed specifically to protect against these types of attacks.

CIFS shares many of the same security issues with NFS. However, most of the issues can be avoided by enabling security features included with the protocol. Man-in-the-middle and session hijacking attacks, in addition to replay and spoofing attacks, can be avoided by message and session authentication. This does not suggest that CIFS can be completely secured, but that care must be taken during the configuration step of implementation to make use of all available security controls.

36.5 ENCRYPTION. Many individuals and organizations focus on encrypting data while it is in motion, transiting networks and the Internet. The use of encryption for data while at rest, or stored, is equally important. As previously mentioned, stored data are compromised by security breaches more often than data in transit. With the use of a sufficiently strong algorithm and sufficiently sized key, data that are stored encrypted can be made unusable for those without the ability to decrypt the data. Even if an attack is successful at gaining full control of copy of the data for brute force attempts at decryption, a proper algorithm and key can prevent the data from being decrypted for a reasonably sufficient period—long enough that any personally identifiable information (PII) or other sensitive, confidential, or proprietary information would be of no value except to historians.

When a system is lost or stolen, an attacker has, essentially, an unlimited amount of time to gain access to the data. If the data are not encrypted, they can simply be read once the userID and password have been determined. If cracking the operating-system access controls is not successful, an unencrypted disk drive can simply be read on a different computer using the file system or, if that's not possible, using

forensic utilities. With encrypted file systems or volume encryption, only a portion of the data is stored encrypted. Data that are not in one of these encrypted locations are vulnerable—including swap spaces and temporary file locations used by the operating system.

See Chapter 7 in this *Handbook* for more details of cryptography; see Chapter 37 for details of the public key cryptosystem.

36.5.1 Recoverability. Key principles of information assurance (IA) include protecting the availability and utility of data. By its nature, encryption can take away these protections in exchange for protecting the authenticity, confidentiality, and integrity of the data while it mitigates the risk of losing possession of the data. When using data encryption it is important to consider the possible need to recover the data in the event that the user or primary key is unable to be located. There are several ways to accomplish this.

Key escrow is one method to facilitate the recovery of encrypted data. By storing the key with a trusted party, a lost key can be removed from escrow and used to decrypt the data. With public key encryption, additional decryption keys (ADKs) can be used with some encryption tools. Corporate ADKs are keys that can be used during the encryption process to automatically encrypt the data with a key that is tightly controlled and only used by designated individuals for the recovery of encrypted data.

36.5.2 File Encryption. The use of encryption on a file-by-file basis is a good method for securing data. With file encryption, a user can pick and choose which files to encrypt. Files containing sensitive data can be encrypted, while nonsensitive files are stored without encryption. This has the least impact on a system, but it puts most responsibility on a user who must determine which files should be encrypted.

Operating system files cannot be encrypted, so configuration files that may contain information on the organization's systems are left as risk—as are application code files. These code files may be for the organization's proprietary application that provides a significant competitive advantage. However, such files are rarely encrypted, as doing so can provide a hindrance for the user who must remember to decrypt each and every file when required.

36.5.3 Volume Encryption and Encrypting File Systems. Both volume encryption and encrypting file systems offer protection to data and can be easier for users than per-file encryption. The ease of use comes from the single operation required to access multiple files. Depending on its configuration, a location can remain decrypted and accessible for a few minutes or several hours.

A critical difference between file encryption and volume encryption is that decryption is carried out at the driver level as data are moved from disk into RAM. The entire volume is *not* decrypted. However, input/output (I/O) to and from the mounted volume can be significantly slower than from an unencrypted disk drive.[8]

In addition, dynamic decryption ensures that there is no extensive plaintext version of the original materials stored on the hard disk or in virtual memory for cryptanalysis in case of an aborted shutdown.

With partial encryption using file encryption or encrypted volumes, only a portion of the data stored on the disk and specifically written to the encrypted partition is protected. Usually, operating-system and application files are not encrypted, so a stolen or otherwise compromised disk is vulnerable to attackers who would gain access to any file not saved into the encrypted volume or file system. If a user forgets to encrypt

a file or folder containing sensitive data, then those unprotected data are accessible to anyone who can access the operating system or read the disk drive using forensic or other disk utilities.

36.5.4 Full Disk Encryption. Author Ryan Groom lists cogent reasons for using full disk encryption on laptops.[9] These can be restated for any system. The primary reasons to use full disk encryption are that it

- Protects data if a drive is lost or stolen,
- Is safer and more effective than volume encryption or encrypting file systems,
- Can be transparent to users, and
- Helps comply with legal and regulatory issues.

Full disk encryption secures the file system and operating-system files but leaves a small boot portion of the drive unencrypted. The unencrypted region allows the encryption software to load, to request the password, passphrase, or token needed to initiate dynamic decryption of the drive contents on demand, and to continue loading the operating system.

Depending on the solution chosen, users can see little difference in functionality or performance between a system using full disk encryption and a system that does not do so. The primary visibility to users is that on a system boot, the user has to identify and authenticate to enable decryption and continue to system boot. System performance is almost undetectably reduced, with only minor delays during the boot of the system while significant amounts of operating-system programs and data from the disk are decrypted, and again on shutdown as unencrypted data are cleaned up to prevent readability without authorization. Although I/O may be slower than on unencrypted systems, most file accesses do not require large volumes of I/O, so overall, these minimal effects on users are greatly outweighed by the protection afforded by full disk encryption.

As with volume encryption, full disk encryption involves dynamic decryption of ciphertext as it flows from disk to memory buffers. With modern data-transfer speeds and processor capabilities, any performance delay due to on-the-fly decryption once the operating system has been loaded is negligible in practice.

With full disk encryption, the entire contents of the drive are protected. Even with full physical access to the disk (e.g., by installing it into another computer under the attacker's control) or with a bit-for-bit copy of the encrypted disk, an attacker must break the encryption in order to gain any information—an almost impossible task with the key sizes currently in use, except perhaps by government cryptanalysis labs using massively parallel architectures for brute-force cracking. With strong encryption, management may be able to satisfy the concerns of clients if an organization has to disclose the loss of equipment and must provide an assurance that, even though the disk was lost, the client and organization's data cannot be accessed.

36.5.5 Vulnerability of Volume, File System, and Full Disk Encryption. As strong as the protection provided by volume, file system, and full disk encryption, there is one significant weakness. Once a user is authorized to access the data, and the operating system dynamically decrypts data as required, the system is vulnerable to attacks by any interloper who has physical access to the unlocked, unprotected session. For example, if a system contains sensitive data and is connected to

a network, any attack over the network can potentially compromise the data when the authorized user's session allows access to the decrypted data.

This vulnerability must be stressed to users who may misunderstand the implications of encryption, especially those who insist on storing sensitive data on their laptops. While full disk encryption protects the data when the system is not booted, once the user decrypts the disk at boot the data are available not just to them but to anyone else who gains access to the system. Systems must use defense-in-depth strategies with personal firewalls to prevent unauthorized network access to the system; users must maintain physical possession of the system, especially once they have entered their decryption key. If a user with a Windows operating system on a laptop locks the screen and then walks away, the data are protected only by the strength of the system password.

Especially sensitive data should be encrypted at the file level. Alternatively, volume and file system encryption can be used with reasonable timeouts applied. Both of these options provide increased protection of data while a system is booted. Combined with full disk encryption, the risks to data are greatly reduced. Some application programs offer options for data encryption (e.g., MS-OUTLOOK includes options for encryption of the PST files containing all the user's data). And what could be worse for an attacker who breaks the full disk encryption only to find that the information is super-encrypted one or more levels, if file encryption, volume encryption, and full disk encryption are all in use?

An important point with the discussion of multilevel encryption is that it requires a strong commitment from the organization to support security at that level. Full disk encryption can be enabled at a global level through group policies or other technology-driven solutions. However, as mentioned above, a user must selectively choose to encrypt sensitive data, placing the responsibility and accountability on them. However, with a multi-layered encryption approach, there is still a layer of protection against a user forgetting or deliberately avoiding encrypting a sensitive piece of data.

Many data privacy laws have arisen over the last several years, including differing laws for the majority of U.S. states and the European Union. Many of these laws require notification of impacted users when data is lost. However, many of these laws include "safe harbor" from notification if the data was encrypted on the lost or stolen device or media. (The authors strongly recommend that you consult with your legal counsel for proper interpretation of the laws for your jurisdiction and/or the jurisdiction of any data subjects potentially impacted by a breach or data loss.)

36.5.6 Database Encryption. For many organizations, the databases that are the primary location for storing data are also an excellent target for attackers. By employing database encryption, the time an attacker needs to gain access can be greatly increased.

Databases can be protected by placing them on a system that can only be accessed via a secured connection and one that employs volume or full disk encryption. In addition, there are database encryption functions or external tools designed to protect the data held within these stores even without disk encryption. These tools employ encryption at the field (individual data element), row (a collection of fields that align in a row when seen in a tabular view), and full database encryption. One of the advantages of database-specific encryption is that it can support strict compliance with legal requirements for data protection. For example, PII such as names, Social Security numbers, medical information, and so on can be protected at the highest level of security available in case unauthorized personnel or intruders gain access to the database.

James C. Foster, in an article for *SearchSecurity.com*, stated that organizations "often jump into database encryption as a quick fix for compliance without considering several key factors. The greatest among these factors is the speed or performance of the application, because poorly implemented database encryption could impact production applications."[10] Foster recommends four "simple guidelines that will help you secure your database without impeding the business you're trying to protect," which are expanded on below:

- Do not encrypt foreign or super keys. These keys are used for indexing, and encrypting them can negatively impact the utility of the database. Since these keys are not encrypted, the keys should never contain information that should be protected—such as using a customer's Social Security number or credit card number as a key to link tables.
- As with any use of encryption, symmetric algorithms are faster than asymmetric keys. However, if all data are encrypted with a single key, this key must be well protected or else an attacker finding the key could, quite literally, have the key to the kingdom of an organization's primary data store.
- "Full database encryption is rarely advised or a feasible option. Security best practices would teach you to encrypt everything with multiple keys and differing algorithms." However, technical staff should evaluate the effects on performance.
- "Encrypt only sensitive data [columns]. This is typically all that is required or recommended by regulations and, after all, is what needs protection."

36.5.6.1 Improving Vendor Provided Options. Database vendors such as Oracle Corporation and Microsoft have encryption options that are specifically designed to protect the information in their databases. These options have improved over time and are becoming more robust and mature.

Microsoft SQL Server 2005 offers enhancements to column encryption. Also introduced was "an integrated and hierarchical infrastructure for managing encryption keys." The product documentation continues, "Built-in encryption functions and application programming interfaces (APIs) make it easier for an organization to create an encryption security framework."[11]

Oracle Database 10*g* Release 2 improves on existing encryption options within Oracle databases by introducing Transparent Data Encryption (TDE). When using TDE, a database administrator is able to specify that a column needs to be encrypted, and the database automatically encrypts data during insert operations and decrypts the data during selects. This can be achieved "without writing a single line of code."[12] Arup Nanda provides a good overview of this feature in the September/October 2005 issue of *Oracle Magazine*.[13]

36.5.6.2 Implementation Considerations. As with any implementation of encryption, careful consideration must be given to determining the method and process for implementing the solution, as well as to the data that are being encrypted. Encrypting a key field needs to be avoided. If sensitive data are in the key field, there may be significant work required to create new key fields and recreate table linkages, or the organization may decide to accept the performance degradation that could occur with encrypting the key field. That is, assuming that encrypting the key field does not cause the database to become unusable.

The costs associated with implementing an encrypted database solution must also be weighed against the business risk. For companies that have little data needing encryption, database encryption may be inappropriate. Legal and regulatory requirements also need to be considered as database encryption may be mandated in order to protect data and to avoid criminal or civil liabilities. The attendant loss of public and customer confidence, in the event of a data breach, is also a powerful incentive.

36.5.7 Smart Phone Encryption. Users may store confidential information on mobile phones and tablets; typical entries include contact entries with phone numbers and sometimes with ancillary sensitive data such as passwords, personally identifiable information (e.g., government-issued identification numbers), and call records. Such devices may also be used as if they were flash drives, with potentially gigabytes of sensitive data downloaded from other sources and carried in a pocket, briefcase, or handbag—and therefore easy to steal or to lose.

Another factor that can be significant for some users is that under U.S. law at the time of writing (in June 2013), a suspect who is questioned, interrogated, or arrested cannot normally be forced to divulge the decryption code.[14]

Phones using Android 2.3.4 or later usually come with integrated total encryption; the process typically takes about an hour, ideally starts with a fully charged battery and connection to a power supply, and must not be interrupted. Interruption of this encryption process can damage or delete the data stored on the phone and requires a factory reset that wipes all current data and personal settings from the device.[15]

Apple iOS and Microsoft Windows Phone 7 also include encryption functions with varying coverage. Third-party software is available for all the operating systems discussed above.[16]

In March 2013, researchers at the Friedrich-Alexander University discovered how to access data encrypted on a version of the Android operating system:

> The team froze phones for an hour as a way to get around the encryption system that protects the data on a phone by scrambling it. … The attack allowed the researchers to get at contact lists, browsing histories, and photos. … [They] put Android phones in a freezer for an hour until the device had cooled to below −10 C. … [Q]uickly connecting and disconnecting the battery of a frozen phone forced the handset into a vulnerable mode. This loophole let them start it up with some custom-built software rather than its onboard Android operating system. The researchers dubbed their custom code *Frost—Forensic Recovery of Scrambled Telephones*.[17]

36.6 DATA DISPOSAL. A final consideration for securing stored data is the disposal of the media that contains the data.

For sanitizing electronic media, United States Department of Defense standard 5220.22-M can provide guidance.[18] Essentially, the media should be sanitized so that the data originally written to it cannot be recovered. One method for achieving this goal can include the following steps:

1. Erase the data
2. Write random or meaningless data to the media
3. Erase the data
4. Repeat until the desired level of sanitization is met

For information stored on paper, the paper should, at a minimum, be shredded before being discarded. Use of a cross-cut shredder is necessary as it increases the difficulty

of piecing the documents back together. In one notorious incident, confetti examined after it was thrown during a Thanksgiving Day parade in New York City in 2012 proved to be easily read horizontal shreds of confidential files from the Nassau County Police Department. "There were whole sentences, license plate numbers, and police reports."[19]

There are additional steps that can be taken if needed, and facilities are available to accomplish these, including burning the shredded paper and mixing it with water to speed its deterioration. Over the past several years, vendors specializing in onsite paper destruction have become commonplace. These vendors bring trucks to a client's site and collect paper that is then shredded onsite before being taken to a facility that recycles, burns, or otherwise disposes of the waste in a secure manner.

For more information about data disposal, see Chapter 57 in this *Handbook*.

36.7 CONCLUDING REMARKS.

The security of stored data is of critical importance. More data breaches occur against data in unsecured storage locations than is compromised during transit. With proper data storage security there is less risk to data from all kinds of threats, internal and external. Using secure channels for writing data to disk and protecting the disks themselves is of increasing importance, as attackers usually have one of two goals—causing systems or data to be made unavailable, or compromising the confidentiality and integrity of data. By combining secure communication channels, data encryption for data at rest, and physical protection of data storage devices, the security of information can be better assured.

36.8 FURTHER READING

Curtis, W. C. *Backup & Recovery*, 2007. Sebastopol, CA: O'Reilly, 2007.

Curtis, W. C. *Using SANs and NAS*, 2002. Sebastopol, CA: O'Reilly, 2002.

Dwivedi, H. *Securing Storage: A Practical Guide to SAN and NAS Security*. Chapter 2: "SANs: Fibre Channel Security." Addison-Wesley, 2006. www.awprofessional.com/content/images/0321349954/samplechapter/Dwivedi_ch02.pdf

Griffin, D. *The Four Pillars of Endpoint Security: Safeguarding Your Network in the Age of Cloud Computing and the Bring-Your-Own-Device Trend*. CreateSpace Independent Publishing Platform, 2013.

EMC Education Services. *Information Storage and Management: Storing, Managing, and Protecting Digital Information in Classic, Virtualized, and Cloud Environments*, 2nd ed. Wiley, 2012.

Hitachi Data Systems. *Storage Concepts: Storing And Managing Digital Data*, Vol. 1. HDS Academy, 2012.

Loftus, A., and C. Kerner, "SAN Lessons Learned," 2003. http://dims.ncsa.uiuc.edu/set/san/src/San_Lessons_Learned.pdf

Loshin, P. *Simple Steps to Data Encryption: A Practical Guide to Secure Computing*. Syngress, 2013.

Shepler, S., B. Callaghan, D. Robinson, R. Thurlow, C. Beame, M. Eisler, and D. Noveck. "RFC 3530 Network File System Version 4 Protocol." April 2003. Proposed Standard Requests for Comment, Internet Engineering Task Force. http://tools.ietf.org/html/rfc3530

Souppaya, M., and K. Scarfone. *Guidelines for Managing the Security of Mobile Devices in the Enterprise*. NIST Special Publication SP800-124 Revision 1. June 2013. http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-124r1.pdf

Systems and Network Analysis Center. "Securing Fibre Channel Storage Area Networks." National Security Agency. February 10, 2009. www.nsa.gov/ia/_files/factsheets/securing_fibre_brochure.pdf

Troppens, U., R. Erkens, and W. Müller, *Storage Networks Explained: Basics and Application of Fibre Channel SAN, NAS iSCSI and InfiniBand*. Wiley, 2004.

Vacca, John. "The Basics of SAN Security, Part I". *Enterprise Storage Forum*. July 23, 2002. www.enterprisestorageforum.com/sans/features/article.php/11188_1431341_2

Verhelst, Wouter. "Securing NFS." Wouter Verhelst's Blog, *Free Software Magazine*. November 26, 2006. www.freesoftwaremagazine.com/blogs/securing_nfs

36.9 NOTES

1. Ponemon Institute, "New Study Reveals Up To 12,000 Laptop Computers Lost Weekly and Up To 600,000 Lost Annually in U.S. Airports," www.ponemon.org/news-2/8
2. M. J. Schwartz, "Lost Laptops Cost $1.8 Billion Per Year: Only One-Third of Missing Laptops Have Full-Disk Encryption for Preventing Data Breaches, Finds Ponemon Study of European Firms," *InformationWeek Security*, April 21, 2011, www.informationweek.com/security/mobile/lost-laptops-cost-18-billion-per-year/229402043
3. D. Klugh, "12,000 Laptops Lost in Airports Every Week," *WMBF News*, October 1, 2012, www.wmbfnews.com/story/19604344/12000-laptops-lost-in-airports-every-week
4. A. Delli Santi, "NJ Audit: Confidential Data on Junked Computers." ABC Action News, WPVI-TV, March 9, 2011, http://abclocal.go.com/wpvi/story?section=news/local&id=8002765
5. M. Kincora, "Strategic Storage: Storage Security—Change Old Habits and Stop Data Theft," SearchStorage.com, November 2005, http://searchdatabackup.techtarget.com/news/1300734/Strategic-Storage-Storage-security-Change-old-habits-and-stop-data-theft
6. W. C. Preston, H. Dwivedi, M. E. Kabay, and S. Gordon, *Storage Security Handbook*, Neoscale Systems, Inc., 2002, www.neoscale.com/English/Downloads/Storage_Security_Handbook/ (URL inactive)
7. Preston et al., *Storage Security Handbook*.
8. For example, observations in June 2013 showed that mounting an entire 7 GB partition encrypted by Symantec PGP Desktop v10.20.0 took ~7 seconds. In comparison, the same 7 GB of data encrypted using WinZip and 256-bit AES encryption took about an hour to decrypt *in toto* on a 7200 rpm two-disk performance RAID 0 set running under 64-bit Windows 7 Professional SP 1 on a 3.2 GHz AMD Phenom™ II X4 955 quad-core processor with 12 GB of RAM. Copying a 1.23 GB file took 85 seconds (~15 MB/sec) on the PGP encrypted volume (located on the RAID 0 hard drive) but only 20 seconds (~63 MB/sec) on the RAID 0 itself.
9. R. Groom, "8 Reasons for Full-Disk Encryption," Business Security, About.com http://bizsecurity.about.com/od/windowsdesktopsecurity/a/top8fulldisk.htm (URL inactive)
10. J. C. Foster, "Look Before Leaping into Database Encryption," Compliance Counselor, SearchSecurity.com, September 29, 2006, http://searchsecurity.techtarget.com/tip/0,289483,sid14_gci1219561,00.html

11. Anonymous, “Improving Data Security by Using SQL Server 2005,” Microsoft® TechNet. http://technet.microsoft.com/en-us/library/bb735261.aspx
12. A. Nanda, “Transparent Data Encryption,” Technology:Security, ORACLE® Technology Network, www.oracle.com/technology/oramag/oracle/05-sep/o55security.html
13. Nanda, “Transparent Data Encryption.”
14. R. Radia, “Why You Should Always Encrypt Your Smartphone,” *ArsTechnica*, January 16, 2011, http://arstechnica.com/gadgets/2011/01/why-you-should-always-encrypt-your-smartphone/
15. How-To Geek, “How to Encrypt Your Android Phone and Why You Might Want To,” 2013, www.howtogeek.com/141953/how-to-encrypt-your-android-phone-and-why-you-might-want-to/
16. UNC-Chapel Hill, “Encrypting Cell Phones,” University of North Carolina at Chapel Hill | Help & Support, March 19, 2013, http://help.unc.edu/help/encrypting-cell-phones
17. BBC, “Frozen Android Phones Give Up Data Secrets: Freezing an Android Phone Can Help Reveal Its Confidential Contents, German Security Researchers Have Found,” *BBC News | Technology*, March 7, 2013, www.bbc.co.uk/news/technology-21697704
18. Anonymous, “DoD 5220.22-M National Industrial Security Program Operating Manual (NISPOM),” 1997, www.usaid.gov/policy/ads/500/d522022 m.pdf
19. C. Boyette, “Sensitive Documents Found in Macy’s Thanksgiving Day Parade Confetti,” *CNN US*, November 26, 2012, www.cnn.com/2012/11/26/us/new-york-confidential-confetti

CHAPTER 37

PKI AND CERTIFICATE AUTHORITIES

Santosh Chokhani, Padgett Peterson, and Steven Lovaas

37.1 INTRODUCTION 37·2
- 37.1.1 Symmetric Key Cryptography 37·2
- 37.1.2 Public Key Cryptosystem 37·2
- 37.1.3 Advantages of Public Key Cryptosystem over Secret Key Cryptosystem 37·3
- 37.1.4 Combination of the Two 37·4

37.2 NEED FOR PUBLIC KEY INFRASTRUCTURE 37·4

37.3 PUBLIC KEY CERTIFICATE 37·5

37.4 ENTERPRISE PUBLIC KEY INFRASTRUCTURE 37·7

37.5 CERTIFICATE POLICY 37·8

37.6 GLOBAL PUBLIC KEY INFRASTRUCTURE 37·9
- 37.6.1 Levels of Trust 37·10
- 37.6.2 Proofing 37·10
- 37.6.3 Trusted Paths 37·10
- 37.6.4 Trust Models 37·11
- 37.6.5 Choosing a Public Key Infrastructure Architecture 37·13
- 37.6.6 Cross-Certification 37·14
- 37.6.7 Public Key Infrastructure Interoperability 37·14

37.7 FORMS OF REVOCATION 37·18
- 37.7.1 Types of Revocation-Notification Mechanisms 37·18
- 37.7.2 Certificate Revocation Lists and Their Variants 37·19
- 37.7.3 Server-Based Revocation Protocols 37·20
- 37.7.4 Summary of Recommendations for Revocation Notification 37·21

37.8 REKEY 37·22

37.9 KEY RECOVERY 37·23

37.10 PRIVILEGE MANAGEMENT 37·25

37.11 TRUSTED ARCHIVAL SERVICES AND TRUSTED TIME STAMPS 37·25

37.12 COST OF PUBLIC KEY INFRASTRUCTURE 37·27

37.13 FURTHER READING 37·27

37.14 NOTES 37·28

37.1 INTRODUCTION. In the 1990s, the use of encryption across the Internet consisted mainly of individuals with Pretty Good Privacy (PGP) and later the Gnu Privacy Guard (GnuPG, originally abbreviated GPG) exchanging secure email and each maintaining a private *web of trust*, today's use of encryption encompasses a much wider range of elements, including proofing, issuance, revocation, identification, federation, bridging, encryption, digital signing, and a myriad of ancillary processes. In fact, what the user experiences is just the tip of the required support structure. Proper management of information involved in an encryption infrastructure (or cryptosystem) is about trust, what is required to establish that trust, and how much to grant.

Back then, the prime use of encryption was outside the network perimeter, where unprotected data sent and received by the organization over public networks were perceived as most vulnerable to disclosure, modification, insertion, deletion, and replay attacks. To protect the data being transported over untrusted networks, the only practical and cost-effective technology is cryptography. Cryptography is at the heart of both virtual private networks (VPNs) and public key infrastructures. For further information on encryption, see Chapter 7. For a review of cryptography, see Chapter 7 in this *Handbook* and *Applied Cryptography* by Bruce Schneier.[1]

37.1.1 Symmetric Key Cryptography. Symmetrical cryptography (also known as *single key* or *secret key*) uses the same key to encrypt cleartext into ciphertext and to decrypt ciphertext back into cleartext. This process is illustrated in Exhibit 37.1.

Although symmetric cryptography (e.g., Advanced Encryption Standard [AES], Data Encryption Standard [DES], Rivest Cipher 4 [RC4]) has specific advantages in that it has a dense keyspace (any integer) and is computationally very fast, it has a fundamental weakness: Both the originator and each recipient must have the key. Key sharing has two management issues:

1. Key exchange must be performed: Somehow each user must have the same key. This exchange is generally accomplished out of band.
2. There is weak confidentiality, control, or authentication: Anyone with the symmetric key could have encrypted the cleartext or decrypt the ciphertext.

37.1.2 Public Key Cryptosystem. In contrast to secret key cryptosystems, public key cryptosystems (PKCs) use pairs of related keys that are generated together. It is accepted that it is infeasible to generate one key from the other. The ciphertext produced by one key can be decrypted only with the other member of the same key

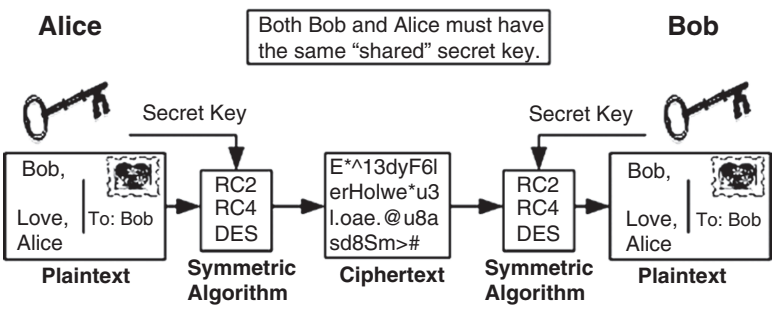


EXHIBIT 37.1 Symmetric (Secret) Key Encryption

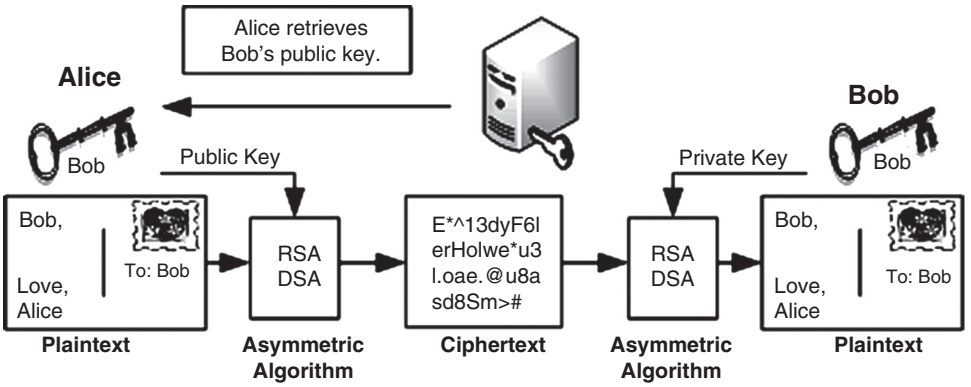


EXHIBIT 37.2 Asymmetric or Public Key Encryption

pair. One of these keys is kept secret (the *private* key) and the other is published for all to use (the *public* key). It does not matter which is designated as which, but once designated, the key cannot be changed. In practice, encryption with one key is often faster than with the other. This one is generally selected as the public key.

In the simplest form, to conceal a message in transit so that only the desired recipient may read it, the cleartext is encrypted using the recipient's public key, as shown in Exhibit 37.2. Only the recipient's secret key can decrypt the transmitted ciphertext. Similarly, to verify message integrity and authenticity, it is possible to encrypt information with a sender's private key. This allows anyone with access to that sender's public key to validate the message by decrypting the ciphertext successfully.

37.1.3 Advantages of Public Key Cryptosystem over Secret Key Cryptosystem. For securing data transmissions, public key cryptosystems are preferred over secret key cryptosystems for these reasons:

- Public key cryptosystems require fewer keys to manage: Each party (n) has a key pair, so the total number of keys is $2n$, instead of being proportional to n^2 as for secret key cryptosystems.
- Because private keys need not be distributed or otherwise managed, public key cryptosystems require only demonstrated integrity and authenticity of the public keys themselves. Users (the *relying parties*) must have assurance that their public keys truly belong to the publishers. This requires signing by a trusted third party or by a mutually trusted source.
- Because no secret keys are transmitted over any network, PKCs are not susceptible to compromise even when public keys have to be changed. PKCs can be used to encrypt temporary keys (*session keys*) that can be used one time for secret key cryptography to obviate the heavier computational load of the PKC.
- To encrypt a message so that multiple PKC users can receive and decipher the ciphertext securely, PKC software can create a session key. This secret key is then encrypted with each recipient's public key separately and sent with the ciphertext to all recipients without compromising confidentiality. Each recipient can then extract the key used to encrypt the file.

EXHIBIT 37.3 Symmetric versus Asymmetric Encryption

Type	Keyspace	Speed
Asymmetric	Sparse	Slow
Symmetric	Dense	Fast

PKC-based digital signatures can also provide the basis for nonrepudiation in the event of a dispute. Only the possessor of a private key could have sent a message decrypted by its public key. In contrast, because of the use of shared secrets, symmetric secret key cryptosystems alone cannot reasonably support nonrepudiation.

37.1.4 Combination of the Two. At this point, it may be easiest to consider the elements of the different types of cryptography, as shown in Exhibit 37.3.

The difference is similar to that between electronic data interchange and e-commerce. Symmetric encryption is good for a very few exchanges with people you trust; asymmetric encryption is good for many, many small exchanges with people or devices you may have never met.

Today, the most common use is to combine both types: For a document to be emailed securely, a symmetric algorithm is selected and a random key generated. The document is encrypted using the symmetric algorithm and key. The symmetric key is then encrypted with the asymmetric public key of each recipient and is then added as a header to the document.

For digital signing, a similar mechanism is used except that a hashing algorithm (e.g., Secure Hash Algorithm [SHA]) is used to create a hash value of the document, and the hash is then encrypted (signed) with the private key of the originator and accompanies the document. The originator's public key can be used to decrypt (verify) the hash, and the hash is used to verify the integrity of the document.

37.2 NEED FOR PUBLIC KEY INFRASTRUCTURE. The PKC depends on the integrity of each public key and of that public key's binding to a specific entity, such as a person, an institution, or a network component. Without mechanisms for ensuring integrity and authenticity, a relying party is vulnerable to masquerading attacks through public key substitution.

To illustrate, suppose that ABC Company wants to send a confidential message to XYZ Corp. that no one else can read. ABC could use XYZ's public key to encrypt the message, although, for the sake of efficiency, ABC probably would use a symmetric algorithm to encrypt the message and the public key algorithm to encrypt the symmetric key. However, if ABC can be tricked into using an attacker's public key as if it were XYZ's public key, then the attacker would be able to decrypt the message. This technique is known as public key *spoofing*.

Such public key spoofing by the attacker would, however, make it impossible for XYZ to read the message from ABC, as was originally intended. Therefore, such an attack probably would continue with the attacker reencrypting ABC's message, using XYZ's real public key, and sending it on to XYZ. Such interception and reencryption is an example of a *man-in-the-middle* attack. However, if the sender also signed the original message, this is something that the attacker could not duplicate, so the message could not be changed but only intercepted.

Another example of breaching the connection between a public key and its owner involves digital signature verification. Suppose ABC wants to verify XYZ's signature. If the attacker could trick ABC into using the attacker's public key as if it were XYZ's public key, then the attacker would be able to sign messages masquerading as XYZ using the attacker's private key. ABC would unknowingly use the replaced public key and be spoofed into thinking that the message actually was signed by XYZ.

This same problem exists with the core of Internet commerce today, Secure Sockets Layer v2 (SSLv2). SSLv3 and Transport Layer Security (TLS) can prevent this, but then every participant must have a certificate trusted by the other party. This is where the concept of a *trust chain* comes in. If a trusted third party has signed XYZ's public key, then the attacker would also have to have a key signed by the same authority as well. Otherwise, a warning would appear that the key was not recognized.

A serious issue is the sheer number of certificate authorities (over 100) built into most browsers, from a variety of countries. The browser trusts all certificates issued by any of these by default, and not all are necessarily worthy of trust.

In summary, both for digital signature and encryption services, the relying party must use the public key of the correct party in order to maintain security. There are various manual, electronic, and hybrid mechanisms for the distribution of public keys in a trusted manner, so that the relying party can be sure to have the correct public keys of the subscribers. These mechanisms for distribution and binding of public keys are known as a Public Key Infrastructure [PKI]).

The beauty of a signed public key is that the trust is inherent with the signature provided the signer is trusted. Unlike a *web of trust,* such as that used by the original PGP which requires all keys be accepted individually by each user, with a *chain of trust* a single signer (the root certificate authority) can provide trust to millions of certificates. Trust need not be absolute; it may be contextual. For example, a certificate signed by an employer can be trusted for elements relating to employment but not for credit cards. See the discussion of trust levels in Section 37.6.1.

37.3 PUBLIC KEY CERTIFICATE. The technique that is most scalable uses a public key certificate issued by a trusted party called the certification authority (CA). A CA issues public key certificates to the various subscribers by putting together subscriber information and signing the information using the CA's private key. The generally accepted standard for public key certificates is the X.509 version 3,[2] as defined in RFC 5280 and subsequent text.

Note: Hierarchial structures may use one or more ***root*** CAs and each Root CA may have multiple ***intermediate*** CAs which issue the end user and server certificates

X.509 certificates are expressed in Abstract Syntax Notation 1 (ASN.1), which is a complex binary notation. In order to be passed through email, certificates are usually MIME (aka Base 64) encoded, expressing the binary syntax in ASCII characters.[3]

The advantage to using a CA and a chain of trust is that by trusting the root, the user automatically trusts all keys that it has issued whether or not the user has ever seen them.

Each CA's certificate may contain this key information:

- Version number of certificate standard
- Certificate serial number (unique for every certificate issued by the CA)
- Algorithm and associated parameters used by CA to sign the certificate
- CA name

- Validity period for the certificate
- Subscriber name
- Subscriber public key, public key algorithm, and associated parameters
- CA unique identifier (optional)
- Subscriber unique identifier (optional)
- Extensions (optional)
- CA's digital signature

The relying parties require the CA's public key so that they can verify the digital signatures on the certificates issued by the CA. The relying party must trust the CA's public key, most likely obtained during the registration process. Once the signatures are verified, relying parties can use the subscriber name and subscriber public key in the certificate with as much confidence in the accuracy of the information as they have in the trustability of the CA.

In some situations, a CA may need to revoke the binding between a subscriber and that subscriber's public key. For example, the subscriber private key may be compromised (i.e., there may be reason to believe that the secret key has fallen into the hands of someone else). Since a public key certificate is an electronic object and can reside in several places at the same time, it is neither practical nor possible to recall, delete, or erase all the copies of the subscriber certificate in a distributed environment. Thus, to invalidate a public key certificate by severing the binding between the subscriber and the subscriber public key, the CA creates a list of invalid certificates. This list is called a *certificate revocation list* (CRL). The relying parties must check that a certificate is not on the CRL prior to using the public key in the certificate. If the certificate is on the CRL, the relying party must not use it. The CA signs the CRL to allow the relying parties to verify the CRL's integrity and authenticity. The key information in the X.509 version 2 CRL is:

- Version number of CRL standard
- Algorithm and associated parameters used by CA to sign the certificate
- CA name
- This CRL issuance time
- Next CRL issuance time (optional)
- List of revoked certificates (listing these items for each certificate):
  - Certificate serial number
  - Time CA was notified of revocation
  - Extensions related to the revoked certificate (optional)
- Extensions related to CRL (optional)
- A's digital signature

It is important that a certificate contain only those elements specifically required for operation and that these be used properly otherwise confusion and improper/unexpected use may result.

Probably the most misused element is the *Key Usage* extension. (It is used so badly that often a second extension, the *Extended Key Usage,* is used to clarify the first.) The most common error is to have the *Non-Repudiation Bit* set on a key intended for

identity only. Nonrepudiation should be asserted only on a key intended as a legal signature, not for identity alone.

A dangerous extension (and one that really does not belong in a user certificate) is the *SMIME Capabilities* extension. This only has meaning in the context of specific email clients but can cripple the use of encryption. Although it may have had some use in early mail applications, today it only establishes a maximum allowed symmetric encryption strength, which may be less than desired. If this extension is not present, RFC 5280 requires a default to triple-DES.

In the case of certificates, less is more. The minimum number of fields and extensions that is required for the assertions the certificate is intended to make is best. PKI systems in the past have failed when too many, and sometimes mission-conflicting extensions, were added.

Particular care needs to be taken that the CA server vendor chosen does not add anything not specified by the certificate template in the Certificate Policy.

37.4 ENTERPRISE PUBLIC KEY INFRASTRUCTURE. The *use* of a certificate is relatively simple, but establishing trust that the certificate is valid and appropriate for use requires a complex set of back-office elements, illustrated in Exhibit 37.4. Each group of users covered by a CA is called a *domain*. Subscribers in a domain receive public key certificates from the appropriate CA. The CA is responsible for generation of subscriber certificates and for CRL generation. The CA posts these signed objects to the repository where the relying parties can obtain them. The CA also archives the certificates and CRLs in case they are required in the future to resolve disputes among the subscribers and the relying parties.

The registration authority (RA) is the trusted representative of the CA and is responsible for authenticating the subscriber identity. The RA typically performs these functions:

- Authenticates (proofs) the subscriber's claimed identity. For example, the RA could require the subscriber to provide a valid photo ID, such as a driver's license or a passport for minimum assurance. Both I-9 authentication, as required by the

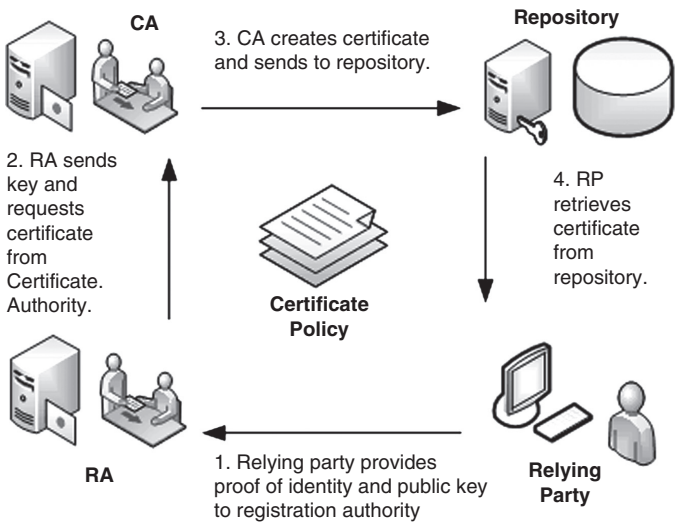


EXHIBIT 37.4 Certificate Issuance Cycle

Immigration Reform and Control Act (IRCA), and a Local Agency Check or a National Agency Check (LAC/NAC) plus "Need to Know" may be required for a higher level.[4]

- Obtains the subscriber public key from the subscriber.
- Provides the CA public key to the subscriber. A trust anchor is a CA's public key that the relying party trusts. This trust generally is established by obtaining the public key from a trusted source using trusted means, such as physical hand-off or via Secure Sockets Layer (SSL) from a trusted or known Website. The CA public key becomes a subscriber trust anchor.
- Sends the certificate creation request to the CA. Typically, the RA creates an electronic mail message containing the subscriber name and the subscriber public key, digitally signs the message, and sends the message to the CA. Other transport means, such as manual or on the Web, also are appropriate as long as there is assurance that the subscriber identity and the public key are not changed. X.509 standard does not specify a protocol for certificate generation requests. The Public Key Infrastructure for X.509 Certificate (PKIX) working group of the Internet Engineering Task Force (IETF) has developed Internet standards in this area.[5]

37.5 CERTIFICATE POLICY. To ensure the security of the PKI, the PKI components need to operate with a high degree of security. To ensure this:

- Private keys must be kept confidential.
- Private keys must be used only by the owners of the keys.
- Trust anchors' public key integrity must be ensured.
- Initial authentication of the subscriber (private key holder and the subject of the public key certificate) must be strong so that identity theft does not occur at the point of certificate creation.
- CA and RA computer systems and applications must be protected from tampering.
- Requirements for level of trust must be clearly defined

The Certificate Policy (CP) must specifically enumerate the certificate contents, both fields and extensions. Anything absent from the CP should not be found in the certificate.

In addition to the security requirements and in order to facilitate electronic commerce, the PKI must address obligations of all parties and their liabilities in case of dispute. These issues of security, liability, level of trust, and obligations are articulated in a CP.

According to the X.509 standard, a CP is "a named set of rules that indicates the applicability of a certificate to a particular community and/or class of application with common security requirements."[6] A certificate user may use a CP to decide whether a certificate, and the binding implied between the certificate and its owner, is sufficiently trustworthy for a particular application. The CP addresses security and obligations of all PKI components, not just the CA; this includes the CA, RA, repository, subscriber, and relying party.

A more detailed description of the practices followed by a CA in issuing and managing certificates is contained in a certification practice statement (CPS) published by or referenced by the CA. According to the American Bar Association's *Digital*

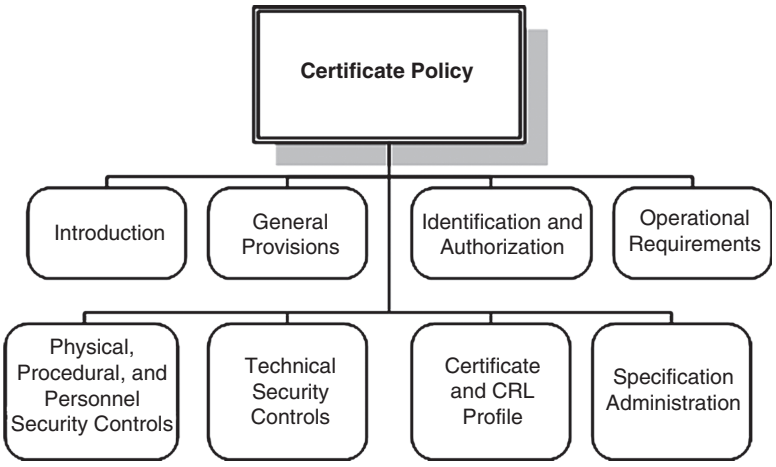


EXHIBIT 37.5 Elements of a Comprehensive Certification Practice Statement

Signature Guidelines (hereinafter referred to as ABA *Guidelines*), "a CPS is a statement of the practices which a certification authority employs in issuing certificates."[7]

Although a CP and a CPS both address the same topics, the CP defines the security requirements and obligations for an enterprise PKI, and the CPS describes how these requirements are satisfied by the enterprise PKI.

The CP and CPS also are used differently. The CP forms the basis for cross-certification across enterprise boundaries to facilitate secure, inter-enterprise electronic commerce. An object identifier (OID) pointing to the CP (which may be a Universal Resource Locator [URL]) is used to create a certificate that can be put into the "Certificate Policies" extension of X.509 certificates. The OID thus enables relying parties to learn the care taken during the generation of certificates, recommended usage, and obligations of the various parties.

Since certificates can be created with different levels of trust and can be either based in software or hardware (more secure), often the "Certificate Policies" extension is the only indicator of the level of trust that should be placed in a certificate.

The CPS enables PKI personnel to use and administer the PKI components. The CPS also forms the basis for compliance audits in order to ensure that the PKI components are operating in accordance with the stipulations of the CPS. Exhibit 37.5 illustrates the components of a comprehensive CPS. Components are divided further into subcomponents, which in turn are divided into elements. Components may be appropriate for various PKI entities but may be applied in the same way or differently. For example, technical security controls may apply to CA, RA, subscribers, and relying parties. These controls may be different for each of these entities, being most stringent for the CA, then the RA, and then the subscribers and relying parties.[8]

A sample policy may be found at www.verisign.com/repository/vtnCp.html

37.6 GLOBAL PUBLIC KEY INFRASTRUCTURE. The principles of an enterprise PKI with a single CA can be extended to support global, secure, electronic commerce by relying on multiple CAs and/or CAs to certify other CAs and each other. How the CAs cross-certify each other is also called *trust model, trust graph*, or *PKI architecture*. For one person to communicate securely with another, there must be a

EXHIBIT 37.6 Trust Level Determination

	Required Trust Level			
Category	**1**	**2**	**3**	**4**
Inconvenience or distress	Low	Med	High	High
Financial loss	Low	Med	Med	High
Harm to agency programs or public interests	N/A	Low	Med	High
Personal safety	N/A	N/A	Low	Med/high
Civil or criminal violations	N/A	Low	Med	High
Information classification				
Confidentiality	Low	Med	High	High
Integrity	Low	Med	High	High

Note: the trust ***level*** is sometimes confused with trust ***factors*** (something you have, something you know, something you know). A Smart Card may be a Level 3 (Medium, hardware) but part of a two factor (username, smartcard) authentication.

trust path from the trust anchor(s) of the relying party to the subscriber whose signature needs to be verified or to whom an encrypted message is to be sent.

37.6.1 Levels of Trust. As referenced in Office of Manpower and Budget Memorandum OMB M04-04, Section 2.1, there are four basic levels of trust[9]:

Level 1. Little or no confidence in the asserted identity's validity

Level 2. Some confidence in the asserted identity's validity

Level 3. High confidence in the asserted identity's validity

Level 4. Very high confidence in the asserted identity's validity

Each level requires a different initial authentication of identity, and higher levels (or those used for classified information) require background investigations. Level 3 is also known as Medium Assurance and is divided into two forms: (1) software (certificates and keys can be exported and moved between devices) and (2) hardware (certificates may be exported but keys and cryptographic functions are performed on a specific hardware device [e.g., a smart card, USB token, or PC-Card device]). Most commercial PKI uses the equivalent of the Medium level of assurance.

Exhibit 37.6 shows one way to determine the required trust level as a function of threat of misuse.[10]

37.6.2 Proofing. Each assurance level has a proofing (also known as *vetting*) requirement that increases with the trust level involved, as shown in Exhibit 37.7. Essentially none is required for level 1 (Basic) while extensive in-person requirements exist for level 4 (High).

37.6.3 Trusted Paths. The trust model can be viewed as a chain, with its tail a certificate-issuing CA and its head the subscriber (i.e., the subject of the certificate). The subscriber can be another CA or an end entity. To ascertain the trustworthiness of a certificate, it is necessary to start with the relying party trust anchor and to follow in the direction of the chain until the subscriber (of interest to the relying party) is reached.

EXHIBIT 37.7 Trust Levels and Proofing

Level	Title	Proofing	Authentication
1	Default	Anonymous allowed.	None
2	Basic	Simple assertion—may be online.	Password
3	Medium (software)	I–9 employment eligibility verification and authorization. Must be in person.	Software certificate
3	Medium (hardware)	I–9 employment eligibility verification and authorization. Must be in person. Biometrics may be captured.	Hardware certificate
4	High	National agency check or local agency check, background investigation, and authorization required. Final proofing must be in person.	Hardware certificate

Global secure communications require that there be a trust path from every subscriber to every other subscriber.

The relying party can start with its trust anchor and verify the certificates issued by the trust anchor. Once that happens, the public keys can be trusted and used to verify the certificates issued by these CAs. This can be done recursively by the relying party until the public key certificate of the subscriber of interest is verified. Then the subscriber public key can be used to verify digital signatures and to perform encryption. Exhibit 37.8 illustrates this pathway. The arrows represent certificates. This is called a *certification path*.

37.6.4 Trust Models. Examples of trust models within PKI that relate to the trust in a certificate and are different from proofing (trust of an identity) include:

- Strict hierarchy
- Hierarchy
- Bridge

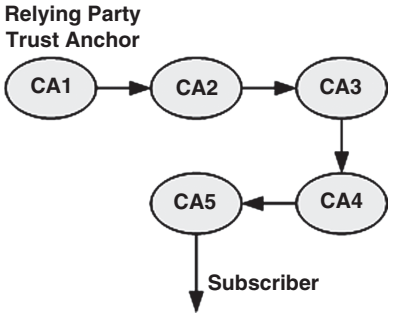


EXHIBIT 37.8 Trust Path through Multiple Trusting Certificate Authorities

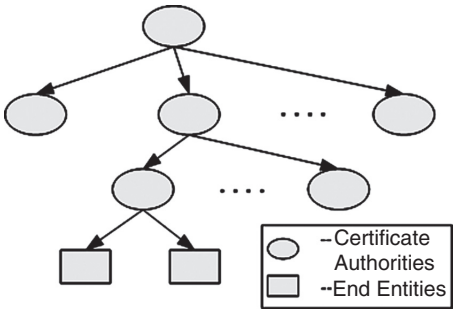


EXHIBIT 37.9 Strict Hierarchical Trust Chain

- Multiple trust anchors
- Mesh (aka anarchy or web)
- Combination

37.6.4.1 Strict Hierarchy. Exhibit 37.9 illustrates a strict hierarchy. It is a tree structure with a single root. In a strict hierarchy, for two parties to communicate with each other securely, they require the public key of their common ancestor as the trust anchor. Verifiable certificate chains require that the parties have a common ancestor. For all parties to communicate securely with each other, they require the single root as the trust anchor, since it is the only common trust anchor.

37.6.4.2 Hierarchy. In a (nonstrict) hierarchy, the subordinate CAs certify their parents. Since the directed graph is bidirectional, any CA can be the trust anchor for the relying parties. But from practical, operational, and performance (i.e., certificate path length) viewpoints, the local CA should be the trust anchor. The local CA is the CA that issued a certificate to the relying party.

37.6.4.3 Bridge. Another trust model is the bridge. Under this model, one CA cross-certifies with each CA from various domains. The domains can be organizations or vertical segments, such as banking or healthcare. Exhibit 37.10 illustrates the bridge CA. The bridge CA is not the trust anchor for any relying party. A CA in the domain of the relying party is the trust anchor.

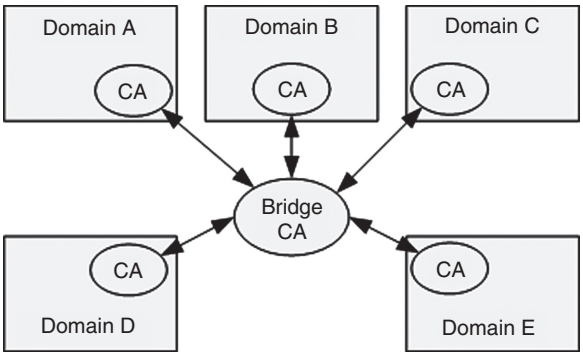


EXHIBIT 37.10 Bridge CA

Within a domain, there are no constraints on the trust model. The domain PKI itself could be organized as any of the trusted models, including bridge, leading to possible layers of bridge CAs.

37.6.4.4 Multiple Trust Anchors. Another alternative is for the relying party to obtain the public keys of the various CAs in a trusted manner and then use those public keys as trust anchors. This approach is attractive when the CAs cannot, or are not willing to, cross-certify, and the relying party needs to communicate securely with the subscribers in the domains of the originating CAs. This approach is called *multiple trust anchors*. Each trust anchor representing a domain could be a single CA or a PKI with a collection of CAs in a trust model.

37.6.4.5 Mesh. The final example of trust model is a mesh (aka *web* or *anarchy*). The term "mesh" describes any depiction representing trust among CAs or certificates without any particular rules or patterns. This model sometimes is known as a *web of trust* and is particularly associated with the original design of Pretty Good Privacy, one of the first popular systems implementing public key certificates. Within this structure, each recipient must explicitly trust each other participant. The major problem is that it does not scale very well beyond a few hundred users.

37.6.5 Choosing a Public Key Infrastructure Architecture. Whether a domain (enterprise) chooses a single CA or multiple CAs for its intradomain operation should be determined from a variety of factors, including:

- Management culture
- Organization politics
- Certification path size
- Subscriber population size
- Subscriber population distribution
- Revocation information

In many situations, the politics or management structure may dictate that there be multiple CAs within the domain. In other words, organizations at business-unit level, regional office level, corporate level, or national level may want to create a CA in order to provide them with a certain degree of control, independence, autonomy, and prestige. How these CAs are organized (bilateral cross-certification, hierarchy, etc.) also will depend on the management and political landscape of the domain. The trust model should be such that it keeps the certification path size manageable; otherwise, end users will see unacceptable performance degradation in obtaining certificates and CRLs and in verifying digital signatures on the certificates and the CRLs.

Similarly, large subscriber populations may require more than a single CA in order to ensure that the CA can manage the subscribers and to keep the CRL size small. If CA products that issue partitioned CRLs are selected, the CRL sizes can be kept manageable even for a very large subscriber population. For further discussion of the CRL issue, see Section 37.7 on revocation alternatives.

When considering interdomain cross-certification, similar issues should be considered.

37.6.6 Cross-Certification. In the simplest form, cross-certification consists of two CAs that certify each other by issuing each other a certificate. The certificates can be stored in specific attributes of the directory entry in a certificate; examples include the cross-certificate attribute pair or the CA certificate.

There are two practical problems with cross-certification. One deals with the commercial products. If the two domains use different products, their CAs may not be able to exchange information to cross-certify, and their directories may not be able to chain to permit the relying parties to retrieve certificates.

The other problem is operational. Before certifying another CA, the certificate-issuing CA needs to make sure that the subject CA is operating in accordance with the acceptable controls, articulated in a CP. The issuing CA asserts the appropriate CP in the "certificate policies" extension of the X.509 version 3 certificate of the subject CA.

In practice, the two CAs cross-certify each other after reviewing each other's CP and after ensuring that the CPs can be claimed to be equivalent. This does not mean that all the security controls and obligations are identical, but they need to offer roughly similar amounts of trust and of obligations and similar liability and financial relief.

When two CAs cross-certify each other, the trust generally is for a limited set of policies through assertions in "certificate policies" extensions, and trust is only bilateral. In other words, trust will not commute; it will remain between the two CAs. The CAs ensure this by inhibiting policy mapping through the "policy constraints" extension. Policy constraint extensions permit differing policy-mapping inhibitions down the certificate chain. In most direct cross-certifications, policy mapping should be inhibited immediately. In the case of cross-certification using the bridge CA model, in order to take advantage of the policy-mapping services of the bridge CA, the policy-mapping inhibition should be different for one certificate (namely the bridge CA certificate).

In addition, the two CAs should use the "name constraints" extension in the X.509 version 3 certificates to ensure that they trust the other domain for the names over which the other has control. The use of this extension also minimizes the chances of name collision.

Exhibits 37.11 and 37.12 illustrate cross-certification examples. These examples are for illustrative purposes only and do not represent real-world entities.

In the case of bilateral cross-certification, policy mapping should be inhibited immediately by using a value of "0" in the "inhibit policy mapping" field of the policy constraints extension in X.509 certificates. When bridge CA is used for interdomain interoperability, a value of "1" should be used in this field. This will permit the issuing CA domain to map its policies to the bridge CA policies and then permit the bridge CA to map its policies to the subject CA domain, in effect mapping from the issuing CA domain to the subject CA domain.

As long as the issuing CA uses its control on inhibit policy mapping, the bridge CA need not use inhibit policy mapping to control the mapping inhibition.

37.6.7 Public Key Infrastructure Interoperability. The complexity of the technology, standards, and products of PKI technology from one domain to another and from one product to another sometimes creates interoperability problems. Yet without interdomain interoperability there can be no global trust, only individual trust.

These factors play a critical role in ensuring PKI interoperability:

- Trust path
- Cryptographic algorithms

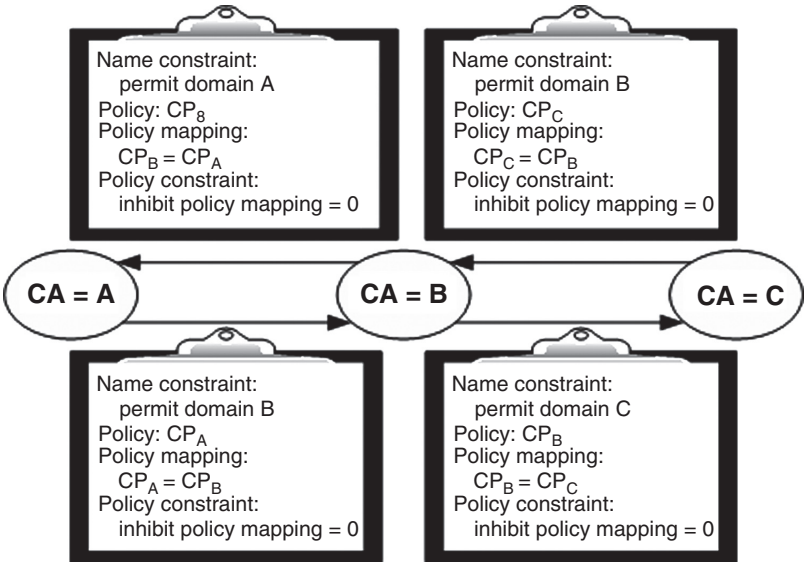


EXHIBIT 37.11 Trust Chain Mapping with Dissimilar Policies

- Certificate and CRL formats
- Certificate and CRL dissemination
- Certificate policies
- Names

37.6.7.1 Trust Path. The communicating parties must be able to form trust paths from their trust anchors to their subscribers. This can be achieved through multiple trust anchors, cross-certification, and other trust models described earlier.

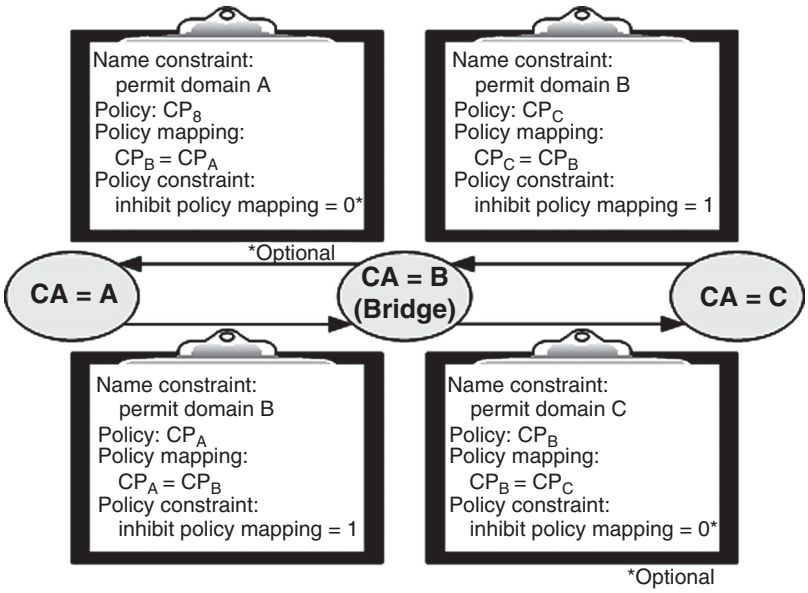


EXHIBIT 37.12 Policy Mapping with a Bridge

37.6.7.2 Cryptographic Algorithms. The communicating parties must implement the cryptographic algorithms, such as hashing, digital signatures, key encryption, and data encryption, used by each other.

In addition, the parties should be able to communicate to each other the algorithms they use. In X.509 certificates and CRL, this information may be contained in the objects themselves, as in the *algorithm* field. In X.509 certificates, for the information being communicated, algorithms such as the digital signature and key encryption algorithm may be carried in the end-entity certificate. The hashing algorithm and the data encryption algorithm can be part of the implicit agreement between the parties or can be carried with the information being communicated. The information also can be obtained from the *supported algorithms* attribute of the X.500 directory entry of the user, although this option is not widely used.

Although the expected public key algorithms used by the CA to create the certificate must be discernible by the recipient (in order to understand the certificate contents), it is important that the certificate not make any assertions as to the symmetric algorithms that may be used. The certificate should be application-agnostic and not impose any rules on applications when not necessary.

In all these situations, the algorithm is identified using the object identifiers. Different organizations may register the same algorithm under their OID arc. Thus, it is important that either the two domains use the same OID for the algorithms, or that their software interpret the multiple OIDs as the same algorithm. For this reason, OID proliferation for algorithms in not recommended.

Variants of the same base algorithm further exacerbate the problems of algorithm interoperability. For example, subtle padding and other differences exist between definitions of the RSA algorithm in the Public Key Cryptography Standards (PKCS) and in the American National Standards Institute (ANSI) X9 committee. Similarly, the Diffie-Hellman algorithm has various modes and various ways to reduce the calculated secret to symmetric key size (i.e., ways to make session keys smaller). Any of these differences in algorithms must be documented through different OIDs so that the OID invokes the appropriate implementation.

It is important that extensions be chosen carefully. Those pertaining to algorithms often have meaning within the context of a particular application. As mentioned earlier, a good example of an extension that should not be placed in a certificate is the *SMIME Capabilities* extension, since it does not set minimum expectations and instead may limit the symmetric key length that can be used by applications.

Such application-specific extensions should be avoided in user/subscriber certificates. The algorithms and key lengths used by the public key (asymmetric) elements are specified in the *Field* values such as *Subject Public Key Information.*

37.6.7.3 Certificate and Certificate Revocation List Format. The communicating parties must share, or must be able to understand, each other's certificate and CRL formats. The most common way to achieve this is to use a common standard, such as X.509. Many times this has not been sufficient due to the ambiguity in the standard and associated encoding schemes, although, over time, those bugs have been worked out. The primary reason today why certificates and CRLs issued by one product may not be understood by another is that either one or both are not compliant with the standard, or one product does not implement all the features of the standard used by the other product.

37.6.7.4 Certificate and Certificate Revocation List Dissemination. The communicating parties must obtain the certificates and CRLs issued by the various CAs in each other's domain. These certificates and CRLs may be obtained from a repository such as an X.500 and Lightweight Directory Access Protocol (LDAP) server. Alternatively, the certificates and CRLs can be carried as part of the communication protocol between the parties, for example, as defined in S/MIME (Secure/Multipurpose Internet Mail Extension) version 3.

The X.500 and LDAP repositories are based on hierarchical databases. Each node in the hierarchical tree structure belongs to an object class. The node's object class determines the attributes that are stored for that node. Examples of attributes are job title, phone number, fax number, and the like. Certificates and CRLs are also attributes.[11]

X.500 and LDAP have defined a standard schema for PKI certificates and CRLs. For certificates, these attributes are *userCertificate*, *cACertificate*, and *crossCertificatePair*. The end-entity certificates should be stored in *userCertificate* attribute. All CA certificates should be stored in the forward element of the *crossCertificatePair* attribute of the subject CA. In addition, all certificates issued to the CAs in the same domain should be stored in the *cACertificate* attribute of the subject CA.

Various revocation lists should be stored in the *cRL*, *aRL*, and *deltaCRL* attributes of the issuing CA as applicable.

If the certificates and CRLs are not stored in these standardized attributes, the relying-party software may not be able to obtain these objects. Furthermore, X.500 directory products still may not always interoperate due to additional complexity of the X.500 standard and to product differences. When implementing X.500 directories and connecting X.500 directory products from different vendors, implementers should allow time to make the products and directories interoperate.

37.6.7.5 Certificate Policies. In order to trust and cross-certify each other, the CAs in two domains need to operate under similar policies. Users in the two domains should be able to accept or reject certificates of each other's domains based on the security requirements of the application, and on the policy under which the certificates were issued.

In order to determine the similarity, or equivalence, of the policies of the two domains, the CP should be written using the IETF standard RFC-2527 framework. The CP is represented using an OID in the certificate. To ensure that the user software accepts and rejects certificates based on the application requirements and on the CP, PKI products should be selected and configured so that the CA asserts the certificate policies, policy mapping, and policy constraints extensions appropriately. The user's PKI-enabling software must process these extensions appropriately and fully in compliance with the requirements of X.509 certificate-path validation rules.

37.6.7.6 Names. The communicating domains must not assign the same name to two different entities. X.500 distinguished names (DNs) are steps in that direction but not sufficient to achieve this.

To illustrate the point, consider, for example, CygnaCom, a company incorporated in the Commonwealth of Virginia. While it is highly unlikely that there is another CygnaCom in Virginia, there is no assurance that there is no CygnaCom incorporated in other U.S. states. Thus, it would be possible that c=US, O=CygnaCom could be asserted by the CAs for several different domains.

In order to avoid this name collision and ambiguity, the *name constraints* extension in X.509 should be used. The CA for one domain can prevent any other entity from using a name registered in that domain. The issuing CA (CA "Y" in this example) uses the *name constraints* extension to assert priority and control over the specified identifier. For example, the first CA that certifies a company called *CygnaCom* in its domain should set the *name constraint* attribute in its certificate for its CygnaCom stating that only its CygnaCom is allowed to issue certificates under the name space c=US, O=CygnaCom. If another CygnaCom were to come along, CA "Y" would ask the second CygnaCom to choose another name in order to avoid name collision. Although this example focuses on the DN, the *name constraint* can be used for any hierarchical name forms, including DN, RFC 822-compliant names, and others.

PKI products should be selected and configured so that the CA asserts the *name constraints* extension appropriately. The user's PKI-enabling software must process this extension appropriately and fully in compliance with the requirements of X.509 certificate-path validation rules.

37.7 FORMS OF REVOCATION. As discussed earlier, a PKI includes mechanisms for key revocation. It is necessary to make provisions for the revocation of compromised keys in order to maintain the trust relationships of any PKI employed in a real-world environment.

The first form of revocation designed was the CRL. It seems the most appropriate form of revocation, given the distributed authentication framework of PKI. The CRL mechanism allows the CA to generate the objects and the relying parties to process them securely without worrying about the security of the servers or system that supply the CRL, and without concerns about the network(s) over which the CRL has traveled.

37.7.1 Types of Revocation-Notification Mechanisms. However, there have been several concerns about the CRL, and these concerns have led to other forms of revocation-notification mechanisms. Many of these mechanisms are variations on the CRL in the sense that these are revocation lists, but they are not complete. The second category of revocation mechanisms defers the processing of revocation information to a server for example, through the Online Certificate Status Protocol (OCSP); see RFC 2560. A third category of mechanisms lets the users check the status of a single certificate from the directory and allows the CA to update the status of that certificate in the directory. A final category lets a CA or another trusted server organize the revocation information in a B-tree.

Which mechanism(s) to choose depends on a variety of factors, such as:

- The communication model (i.e., which class of users is communicating with which other class). For example, if a user communicates with several users who are subscribers to the same CA, a single CRL from that CA will provide relevant information about all those targeted users. If a user is communicating with users who belong to different CAs, each CRL provides information about only one user.
- The directory architecture: Where they are located and what portions of the directory information is replicated or shadowed?
- The communication bandwidth available.
- The bind time (i.e., the time to set a connection with the repository in order to perform retrievals and updates) to access the repository.
- The size of the revocation response from the repository (e.g., the CRL size).

- The processing load on the repository, especially for digital signature *generation* on the revocation information.
- The processing load on the user workstation, especially for digital signature *verification* on the revocation information.

37.7.2 Certificate Revocation Lists and Their Variants. The first set of mechanisms, CRL and its various forms, is the most versatile, effective, and recommended approach for revocation notification. Like X.509 certificates, CRLs are expressed in ASN.1 format. There are several basic types of CRL, and they should be carefully considered, based on the user communication model and anticipated revocation rate:

- Full and complete CRL
- Authority revocation list (ARL)
- Distribution-point CRL
- Delta CRL

37.7.2.1 Full and Complete CRL. The full and complete CRL is a CRL that contains the revocation information for all certificates issued by a CA. This type of CRL is rarely seen; instead a normal CRL includes only information about revoked certificates and not currently valid ones. Expired certificates are not included, and revoked certificates will drop off the CRL when they expire.

37.7.2.2 Authority Revocation List. The ARL is a CRL that contains the revocation information for all the CA certificates issued by a CA; that is, the ARL is a subset of CRL for certificates issued to the CAs only. The ARL is a very desirable mechanism for these reasons:

- It is likely to be short. A CA is likely to certify fewer CAs than other types of subscribers. Also, given that CAs are expected to operate with a great deal of vigilance, and given that CAs are not going to be revoked for reasons such as name change or organizational affiliation change, CAs will be revoked far less often than the end entities. These factors will contribute to making the ARL very small.
- For all of the certificates except one, only the ARL needs to be checked, since in a certificate path all but the last certificate is issued to a CA.

Due to a security flaw in X.509 version 1, a CA should never issue ARLs defined using that version. In X.509 version 1, there is no difference between the CRL format and the ARL format. Since both CRLs and ARLs are signed by the same CA, if an adversary (directory or network adversary) were to supply an ARL to the relying party in lieu of a full CRL, the relying party would have no way of knowing that it had received an ARL instead of the requested CRL. The ARL would not have end-entity revocation information and therefore could mislead the relying party into using the revoked certificate of an end entity.

The X.509 version 2 ARL fixes this security flaw using an *issuing distribution point* extension. An ARL must use this extension and assert a field that states that the list contains only CA certificates. The presence of this field in the signed ARL tells the

relying party that it is not a full CRL. Now, if an adversary were to supply an ARL in lieu of a CRL, the relying party would detect this substitution by using the *issuing distribution point* field.

This is one of the several security reasons that PKI-enabling software must be able to process the various extensions properly in accordance with the requirements stated in X.509 standard.

37.7.2.3 Distribution-Point CRL. Distribution-point CRL is a mechanism that has several useful functions:

- To replicate a CRL
- To consolidate revocation information from the various CAs so that the relying parties need to obtain only one CRL
- To partition the revocation information for the subscribers of a CA into multiple smaller pieces

This latter function, partition, is achieved by asserting the *CRL Distribution Point* extension in the certificate that points to the name entry under which revocation information for the certificate will appear. The partitioned CRL will assert the same name in the *Distribution Point* field of the *issuing distribution point* extension in the CRL.

Since all the partitioned (distribution point) CRLs are signed by the same CA, it is not sufficient for the relying party simply to validate the CA's signature on the Distribution Point CRL. The relying party must match the Distribution Point name in the *issuing distribution point* extension of the CRL with the Distribution Point name in the *distribution point* extension in the certificate.

37.7.2.4 Delta Certificate Revocation List. Yet another way to reduce the size of the CRL is to publish changes to the revocation information since the last CRL. The CRL that contains changes only is called the *delta CRL*, and the CRL to which changes are published is called the *base CRL*. The delta CRL can be applied to any of these CRLs: CRL, ARL, and Distribution Point CRL. In order to construct current revocation information, the latest delta CRL and its base must be used. There is an algorithm that could be used to allow a subset of changes to be applied to an earlier CRL that would still match the digital signature of the new CRL.

37.7.3 Server-Based Revocation Protocols. Server-based revocation uses protocols, such as On-Line Certificate Status Protocol (OCSP) and Simple Certificate Validation Protocol (SCVP). In general, these protocols suffer from several flaws, including these:

- Since the revocation information is produced at the server, the communication channel between the relying party and the server must be secured, most likely by using digital signatures.
- Signed operations will limit server scalability since digital signature generation is computationally intensive.
- Since the revocation information is produced at the server, the scheme requires a trusted server as opposed to an untrusted repository.

- Revocation of a server public key requires a method for checking the server public key status. This method is likely to use the server public key as an additional trust anchor or to rely on a CRL mechanism.
- There needs to be a nonsuppressible mechanism for the CA to provide revocation information to the trusted server; that is, the CA should know whether the revocation information has or has not reached the trusted server. Although a CA itself can act as a trusted server, this is not recommended for security reasons; in addition, we do not want to impose the high-performance requirement on the CA architecture. The trusted server must be a high-performance system.
- There are no standards in the area of CA to provide nonsuppressible mechanisms for transmitting the revocation information to the trusted server.

These mechanisms may be desirable under one of four situations:

1. Need to have thinnest possible PKI client
2. Need to generate revenue for CA services
3. Need to check changing credentials, such as available credit
4. Need to update dynamic credentials, such as the remaining credit line

The last two situations permit the trusted server to provide the revocation information and to check or change the credentials of the subscriber.

Delta CRLs and server-based revocation/authentication protocols such as OCSP (RFC 2560) are standards-compliant and can provide the same information as in the CRL for a single certificate in a significantly smaller bandwidth. They do require some form of acceptable authentication since the original CA will not be available for signing.

37.7.4 Summary of Recommendations for Revocation Notification. The most scalable and versatile revocation-notification mechanism can be achieved by using a combination of:

- CRLs.
- Replication of the CA directory entry, at locations determined by the enterprise network topology, for fast access to CRL.
- Use of ARLs.
- Consolidation of ARLs for all CAs in a domain through the use of distribution points. Consolidation is achieved by placing the name of a CA that can revoke a certificate in the certificate's CRL *Distribution Point* extension.
- Consolidation of all the reason-codes of key compromise for all certificates in a domain through the use of the *Distribution Point* extension. This CRL can be issued very frequently to meet the freshness requirements of the domain. This mechanism makes the CRL mechanism as current as the OCSP.
- Partitioning routine revocation information using Distribution Point CRLs if CRLs become too large.

Several other techniques can help improve CRL retrieval efficiency:

- Repositories may store both enciphered CRLs to send to the relying parties and also deciphered (plaintext) CRLs to perform fast searches. Storing both forms reduces the overhead that would result from using encryption or decryption at the time of each request.
- If the repository does not store any private information, bind operations for retrieval can be configured to require no authentication, thus eliminating another potential performance bottleneck.
- CRL size can be reduced by having a short validity period for the certificates, by using a coarse domain name so that reorganization does not invalidate a name, and by allowing some changes (e.g., name change or transfer) without forcing revocation.

37.8 REKEY. The public key certificates for the subscribers have a defined validity period. Once the validity period expires, subscribers require new public key certificates. There are two primary reasons why public key certificates have a limited life. One relates to the life of a private key based on the potential cryptanalysis threat. Another reason is to help control CRL size since no certificate gets off a CRL until it expires.

No public key should be used longer than the estimated time for brute-force cryptanalysis using current technology (its *cryptanalysis threat period*). At that point, the certificate should be assigned a new public key (i.e., it should be *rekeyed*). However, before the cryptanalysis threat period expires, the same key can be renewed or *recertified*. Certificates can be renewed easily by having subscribers send a digitally signed request to the CA or by having the CA perform automatic renewal. During renewal, any information (other than the subscriber public key) may be changed.

For the foreseeable future, the life of a 1024-bit RSA key can be expected to be 25 years. By reusing the same key with new certificates for as long as possible, the number of past keys required to be escrowed in order to retrieve or validate older files is reduced. The level of trust should also be considered when rekeying, but keys stored in tamper-resistant hardware can be expected to achieve their full lifetime.

Note: NIST SP800-78-3 recommends that after 2013 only 2048 bit keys be used.

Elements using soft keys are inherently at risk to a determined attack and should be used only for lower levels of trust.

Certificates also can be rekeyed easily by having the subscriber send a digitally signed *rekey request message* that also contains the new public key. The message is signed using the current private key so that it can be verified using the current public key. If the subscriber being rekeyed is a CA, these requirements also come into play:

- The relying parties should be able to verify certificate chains after the CA is rekeyed.
- The relying parties should be able to verify CRLs issued by the CA.
- The rekey should not have a ripple effect on the PKI. Just because one CA rekeys, other CAs or end entities should not have to rekey.
- The length of certificate paths should be minimized.
- The operational impact on the PKI entities should be minimal.

A good way to meet these requirements is for the CA to:

- Issue all current valid certificates when it rekeys, without changing the validity periods in the subscriber certificates.
- Continue to sign CRLs with all current valid private keys. This will result in multiple CRLs, all with the same information. A CA private key is considered valid until all certificates signed using that key have expired.

If the CA is a trust anchor, it can use one of two approaches to rekey itself in-band, over the untrusted network:

1. The CA can send out a rekey message that contains its new public key and is signed using the current key. The CA needs to ensure that all the subscribers receive and process the rekey message prior to expiration of the current key.
2. The CA can provide the hash of the next public key and parameters (if the cryptographic algorithm has parameters; RSA does not have parameters, but Digital Signature Standard [DSS] does) with the current key. When it is time to publish the new public key, the CA can publish a new self-signed public key certificate that contains the new public key and parameters as well as the hash of the next public key and parameters.

37.9 KEY RECOVERY. Subscriber public keys can be used to encrypt data-encryption keys (for symmetric key encryption). Such data-encryption keys are used to encrypt data quickly with the lower overhead of symmetric key encryption. Subscribers require their private keys to decrypt the data-encryption keys and thus allow decryption of the data.

It is critically important to distinguish between *signing* keys and *data-encryption* keys. The former (any key that asserts ***nonrepudiation***) may *never* be subjected to key recovery; the latter *may* be protected using key-recovery techniques.

High levels of trust keys may require a separate identification key that may not be used to sign a document; this login key should have only the signing bit and not the nonrepudiation bit set in the Key Usage extension.

Sometimes a subscriber's private key token (e.g., a diskette, hard drive, smart card, etc.) may be corrupted or the subscriber may forget the password associated with the token. Similarly, sometimes a subscriber may not be available, yet the subscriber's employer may need to decrypt corporate information encrypted by the missing subscriber. Key-recovery techniques are designed to meet these emergency needs for access to encrypted information. Inherently, they provide a form of back door to the keys, but they also impose additional overhead costs. Thus, the need to provide key recovery should be balanced carefully against potential costs and complexity.

The two most popular forms of key-recovery mechanisms are:

1. **Key escrow.** Under this form, the subscriber's long-term private decryption key is provided to a trusted third party called a *key recovery agent* (KRA)
2. **Key encapsulation.** Under this form, the subscriber encrypts the data-encrypting key using the public key of the KRA so that the KRA can decrypt the data.

Of these two schemes, key escrow is becoming more widely available in the PKI products because it is simpler to implement at the infrastructure level. It is also

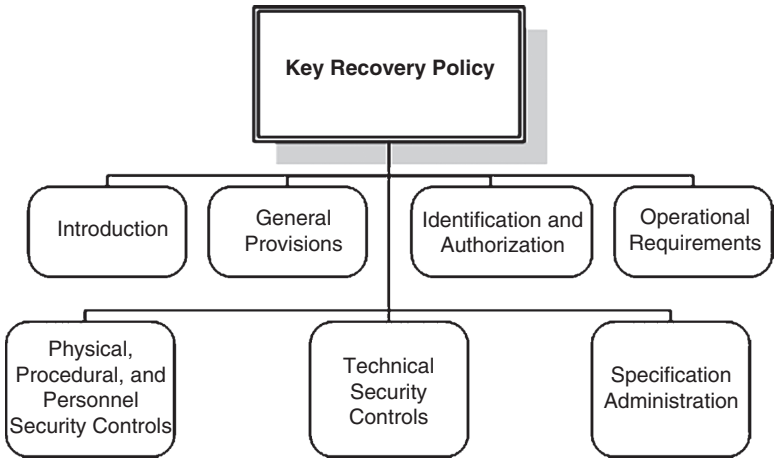


EXHIBIT 37.13 Elements of a Key Recovery Policy

independent of organization boundaries between the sender and receiver of encrypted communications. If a party's private date-encryption key is escrowed, than communications to the party can be decrypted.

Subscribers may always recover their own data-encryption key from the key-recovery system. Authorized third parties, such as a subscriber's employer, also may request keys. Such an authorized party is called a *key recovery requester* (KRR). All of the components are governed by a *key recovery policy* (KRP) and associated *key recovery practices statement* (KRPS). The KRP and KRPS are akin to the *Certificate Policy* and *Certification Practice Statement* but have some differences. One of the main differences is in the technical security-controls sections. There are several requirements to check the communication protocols among the components to ensure confidentiality, integrity, and authorization. Exhibit 37.13 illustrates the components of a key-recovery system, as would be expressed in the KPRS.

A general criticism of key recovery is that it provides secrets to a single party, namely a KRA. One way to mitigate that concern is to share the secret among multiple recipients in a way that requires cooperation (if it is authorized) or collusion (if it is not) among two or more holders of the escrowed secret. For example, superencryption (encryption of ciphertext) can make unauthorized discovery of a key more difficult. The secret key S is encrypted with one recipient's public key (say, K_1), producing a ciphertext, represented as E(S, K_1), and then that ciphertext is super-encrypted using a second recipient's public key, K_2, to produce the ciphertext E(E(S, K_1), K_2). Unlike encryption of the same message for two recipients, where each recipient can decrypt the ciphertext independently, superencryption requires decryption by each recipient in the reverse order of priority. Thus, if a user encrypts a secret key using A's public key and then superencrypts the ciphertext using B's public key, key recovery requires decryption by B using the corresponding private key and then decryption of the resulting ciphertext by A with that user's private key.

To use superencryption so that fewer than all recipients may decrypt the key, the sender encrypts to a first group of recipients and then superencrypts to a second group of recipients. Thus, any one of the members of the first group and any member of the second group of recipients can cooperate to decrypt the secret key.

This technique allows key escrow even in the absence of a formal PKI, such as in informal webs of trust using PGP.

Another solution to making collusion more difficult in key escrow is to split the key using Shamir's *n* out of *m* rule.[12] In a properly implemented key-splitting scheme, parts of a secret key are distributed to *m* destinations and at least *n* recipients are required to reconstitute the secret key. Thus, *n*–1 or fewer persons colluding together cannot determine even a single bit of the escrowed key. Successful collusion requires at least *n* individuals. The split-key approach can be applied to key escrow, in which case the private key can be split and different splits can be provided to different KRAs. Alternatively, the split-key approach may be applied to encapsulation, where the session key can be split, and different splits can be encrypted using public keys of different KRAs.

This strategy also theoretically allows full reconstitution of the secret key to be performed by the authorized recipient of the partial keys, instead of by any of the escrow agents.

37.10 PRIVILEGE MANAGEMENT. The primary purpose of PKI is to provide entity authentication in a global, distributed environment. In most systems and applications, authentication becomes the basis for access control. There are three fundamental ways to implement access control:

1. The systems and applications can perform access control on their own. The PKI continues to provide the authentication framework.
2. The privileges, attributes, roles, rights, and authorizations can be carried in a public key certificate in the *subject directory attribute* extension.
3. The privileges, attributes, roles, rights, and authorizations can be carried in an *attribute certificate*. The X.509 standard is being revised to include the concept of attribute certificates. The attribute certificates carry privileges and authorizations instead of the public key, and thus provide a distributed authorization framework.

User login, signing, and encryption certificates should never be used to carry authorizations or application information. Where needed, separate certificates can be issued for those tasks. Such authorization certificates can be issued daily, or even for shorter periods, eliminating the need for a CRL since the certificate would expire before one would be issued. These authorization certificates could be used in place of Kerberos (the Greek mythological three-headed dog that guards the entrance to Hades) Tickets. Exhibit 37.14 summarizes the pros and cons of each of the three approaches.

These factors should be considered when architecting an access infrastructure. Although the attribute certificate seems to be the latest fad in the PKI world, users should carefully study putting privileges in public key certificates to save the cost of implementing a Privilege Management Infrastructure (PMI) over and above the cost of PKI.

37.11 TRUSTED ARCHIVAL SERVICES AND TRUSTED TIME STAMPS. PKI technology supports global electronic commerce through the use of digital signature technology. Digital signature technology is a detection or passive mechanism. In other words, the technology does not *prevent* someone from modifying data in storage or in transit, nor from impersonating someone else. The technology merely detects that an attempt had been made to modify the data or that someone had tried to impersonate someone else.

EXHIBIT 37.14 Privilege Management

Alternative	Pros	Cons
Application-based access control	Easy to implement. Does not require additional infrastructure, so saves cost.	Need to manage privileges on an application-by-application basis. Synchronization of privileges may be hard as applications increase and as they are distributed. Security may be compromised if privileges are not removed from all applications. Higher operational costs.
Public key Certificate	Easy to add to PKI. Privileges can be managed easily by revoking certificate.	Changes in privileges require revocation of identity certificate. Sometimes this is a small price to pay for savings that result from not having to deploy and operate a separate privilege management infrastructure (PMI). Parties issuing identity certificate may not have authority to bestow privileges.
Attribute Certificate	Privileges can be managed easily by revoking attribute certificates. Change in privilege does not require revocation of public key certificate.	Cost of privilege management infrastructure (PMI).

In a court of law, digital signatures may be disputed long after they were applied, if, for example, the cryptanalysis threat period for the keys has expired. In those circumstances, producing a document with a verified digital signature may not prevent repudiation. A party could claim that the cryptanalysis threat period has passed and the private key might have been discovered or broken by an adversary.

To mitigate both of these threats—data corruption and expiry of the cryptanalysis threat period—trusted archival services are required for transactions with a potential for this kind of dispute. Such an archival service also would be able to safeguard associated certificates and CRLs. Trusted archival services should depend on controls such as physical security, stable media (e.g., write-once read-many [WORM] devices), and appropriate techniques to maintain readability despite changing technologies. Such services should be capable of error-free transcription of data on older media and of translating outdated encoding into more modern media, and into current encoding schemes. For example, Extended Binary Coded Decimal Interchange Code (EBCDIC)–encoded data on nine-track magnetic tapes could be copied onto optical disks using ASCII encoding.

A related technology is the trusted time stamp, where a trusted third party attaches the current valid time to a document and signs it to prove the existence of the document

at a given time. If the document owner does not want to reveal the contents of the document to the time stamp server, then a hash of the document may be stamped and signed instead.

In most applications, a trusted archival service may obviate the need for trusted time stamp service because, for the long term, the trusted archival service can attest to the time of the transaction. For the short term, both parties can date and digitally sign the transaction when it is consummated, rendering it valid as a contract. If the date is not acceptable to any party—it is either too far in the future or in the past—that party can either reject the transaction or invoke some dispute-resolution procedure immediately.

37.12 COST OF PUBLIC KEY INFRASTRUCTURE. One of the misconceptions about Public Key Infrastructure technology is that it is too costly, but these costs should be compared to the alternatives.

Short of taking a major risk, there is no practical technology other than cryptography to protect data in transit over untrusted network. The only choice, then, is between symmetric key cryptography and public key cryptography. Aside from the difficulties of distributing and supporting symmetric, secret keys, such cryptosystems require approximately n^2 keys for a group of n individuals who must communicate with each other.

A secret key cryptosystem requires n^2 keys to be kept confidential, with their integrity maintained; PKI requires managing only n keys. Clearly, it should be cheaper to maintain the integrity of n keys than to manage n^2 keys and their confidentiality and integrity. PKI is needed to manage the public keys.

PKI costs seem large when they must provide global trust and interoperability, something not asked of most systems or infrastructures. Currently, no other technology works as well or is as cost effective as PKI to achieve global, secure, trusted communication and electronic commerce. The alternative is to assume the risk without PKI, or to continue with the last century's approach of paper-based trust.

37.13 FURTHER READING

Adams, C., S. Lloyd, and S. Kent. *Understanding the Public-Key Infrastructure*, 2nd ed. Upper Saddle River, NJ: Addison-Wesley Professional, 2002.

Atzeni, A. S., and A. Lioy, eds. *Public Key Infrastructure, Third European PKI Workshop: Theory and Practice, EuroPKI 2006, Turin, Italy, June 19–20, 2006. Proceedings*. New York: Springer, 2006.

Austin, T. *PKI: A Wiley Technical Brief*. New York: John Wiley & Sons, 2000.

Australian Government Public Key Infrastructure; see: www.govonline.gov.au/projects/publickey/index.asp

Ballad, B. *Access Control, Authentication, and Public Key Infrastructure*. Jones & Bartlett Learning, 2010

Barker, E., W. C. Barker, and A. Lee. *Guideline for Implementing Cryptography in the Federal Government,* 2nd ed., SP 800-21. NIST, 2005. http://csrc.nist.gov/publications/nistpubs/800-21-1/sp800-21-1_Dec2005.pdf (URL inactive)

Barker, E., and A. Roginsky. *Recommendation for Cryptographic Key Generation*. SP 800-113. NIST, 2012. http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-133.pdf

CertiPath. CertiPath X.509 Certificate Policy, version 3.22. CertiPath, 2013. https://www.certipath.com/images/stories/data/policy-docs/2013-05-22%20CertiPath%20CPv.3.22.pdf

Chadwick, D., and G. Zhao, eds. *Public Key Infrastructure. Second European PKI Workshop: Research and Applications, EuroPKI 2005, Canterbury, UK, June 30–July 1, 2005. Revised Selected Papers*. New York: Springer, 2005.

Desmedt, Y. G., ed. *Secure Public Key Infrastructure: Standards, PGP and Beyond*, 2nd ed. New York: Springer, 2013.

ECA Certificate Policies: http://iase.disa.mil/pki/eca/documents.html

Ford, W., and M. S. Baum. *Secure Electronic Commerce: Building the Infrastructure for Digital Signatures and Encryption*. Upper Saddle River, NJ: Prentice-Hall, 1997.

Housley, R., and T. Polk. *Planning for PKI: Best Practices Guide for Deploying Public Key Infrastructure*. New York: John Wiley & Sons, 2001.

Howes, T., M. C. Smith, and G. S. Good. *Understanding and Deploying LDAP Directory Services*, 2nd ed. New York: Macmillan, 2003.

IETF. Public-Key Infrastructure (X.509) (pkix) Working Group Charter, version 4.50. IETF, 2013. http://datatracker.ietf.org/wg/pkix/charter/

Karamanian, A., S. Tenneti, and F. Dessart. *PKI Uncovered: Certificate-Based Security Solutions for Next-Generation Networks*. Cisco Press, 2011.

Kuhn, D. R., V. C. Hu, W. T. Polk, S.-J. Chang. *Introduction to Public Key Technology and the Federal PKI Infrastructure.* NIST SP800-32, 2001. http://csrc.nist.gov/publications/nistpubs/800-32/sp800-32.pdf

Lopez, J., P. Samarati, and J. L. Ferrer, eds. *Public Key Infrastructure: Fourth European PKI Workshop: Theory and Practice, EuroPKI 2007, Palma de Mallorca, Spain, June 28–30, 2007. Proceedings*. New York: Springer, 2007.

Menezes, A., P. van Oorschot, and S. Vanstone. *Handbook of Applied Cryptography*. Boca Raton, FL: CRC Press, 1996.

NIST. *DRAFT A Profile for U. S. Federal Cryptographic Key Management Systems (CKMS).* SP 800-152. NIST, 2012. http://csrc.nist.gov/publications/drafts/800-152/draft-sp-800-152.pdf

Sabo, J. T., and Y. A. Dzambasow. "PKI Policy in the Business Environment," *SC Magazine Asia* (June 2002), http://scmagazine.com/asia/news/article/419807/pkipolicy-business-environment

US Department of Defense External Certification Authority X.509 Certificate Policy, version 4.3. DoD, 2012. http://iase.disa.mil/pki/eca/downloads/pdf/eca_cp_v4-3_final_signed.pdf (URL inactive)

Vacca, J. R. *Public Key Infrastructure: Building Trusted Applications and Web Services*. Boca Raton, FL: Auerbach, 2004.

37.14 NOTES

1. B. Schneier, Applied Cryptography: Protocols, Algorithms, and Source Code in C, 2nd ed. (New York: John Wiley & Sons, 1995).
2. Additional details on X.509 certificates can be obtained from X.509 standard ISO/IEC 9594-8.
3. For more information on ASN.1, see http://asn1.elibel.tm.fr
4. The I-9 "Employment Eligibility Verification" form, properly known as OMB No. 1615-0047, can be found at: www.uscis.gov/files/form/i-9.pdf
5. These standards can be obtained from www.ietf.org
6. RFC 2527, Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework: www.ietf.org/rfc/rfc2527.txt

7. American Bar Association, "Digital Signatures Guidelines: Legal Infrastructure for Certification Authorities and Electronic Commerce." Draft, 1995. www.abanet.org/scitech/ec/isc/dsgfree.html (URL inactive)
8. Further details on CP and CPS format and contents can be found in S. Chokhani and W. Ford, Certificate Policy and Certification Practices Framework, RFC 2527 (April 1998), www.ietf.org/rfc/rfc2527.txt
9. www.whitehouse.gov/omb/memoranda/fy04/m04-04.pdf
10. New York Office for Technology, "Potential Impacts of Authentication Errors," www.oft.state.ny.us/policy/G07-001/table1.htm (URL inactive)
11. For further discussion on repositories, see T. Howes et al., Understanding and Deploying LDAP Directory Services (New York: Macmillan, 1998).
12. A. Shamir, "How to Share a Secret," Communications of ACM 22, No. 11 (1979): 612–613, and W. T. Polk, D. F. Dodson, W. E. Burr, H. Ferraiolo, and D. Cooper, "Cryptographic Algorithms and Key Sizes for Personal Identity Verification," NIST Special Publication 800-78-3, December, 2010. http://csrc.nist.gov/publications/nistpubs/800-78-3/sp800-78-3.pdf

CHAPTER 38

WRITING SECURE CODE

Lester E. Nichols, M. E. Kabay, and Timothy Braithwaite

38.1 INTRODUCTION 38·1

38.2 POLICY AND MANAGEMENT ISSUES 38·1

38.2.1 Software Total Quality Management 38·2
38.2.2 Due Diligence 38·4
38.2.3 Regulatory and Compliance Considerations 38·5

38.3 TECHNICAL AND PROCEDURAL ISSUES 38·5

38.3.1 Requirements Analysis 38·5
38.3.2 Design 38·6
38.3.3 Operating System 38·6
38.3.4 Best Practices and Guidelines 38·7
38.3.5 Languages 38·9

38.4 TYPES OF SOFTWARE ERRORS 38·10

38.4.1 Internal Design or Implementation Errors 38·10

38.5 ASSURANCE TOOLS AND TECHNIQUES 38·15

38.5.1 Education Resources 38·15
38.5.2 Code Examination and Application Testing Techniques 38·16
38.5.3 Standards and Best Practices 38·20

38.6 CONCLUDING REMARKS 38·20

38.7 FURTHER READING 38·21

38.8 NOTES 38·22

38.1 INTRODUCTION. The topic of secure coding cannot be adequately addressed in a single chapter. Unfortunately, programs are inherently difficult to secure because of the large number of ways that execution can traverse the code as a result of different input sequences and data values. The additional pressures to get the applications to market can also take a toll on the secure coding process.

This chapter provides a starting point and additional resources for security professionals, system architects, and developers to build a successful and secure development methodology. Writing secure code takes coordination and cooperation of various functional areas within an organization, and may require fundamental changes in the way software development currently is designed, written, tested, and implemented, as well as a change in corporate culture as it pertains to software development.

38.2 POLICY AND MANAGEMENT ISSUES. There are countless security hurdles facing those writing code and developing software. Today dependence on the

reliability and security of the automated system is nearly total. For an increasing number of organizations, distributed information processes, implemented via networked environments, have become the critical operating element of their business. Not only must the processing system work when needed, but also the information processed must retain its integrity so that it can be trusted in use. Because of a general lack of basic IT organizational discipline, and adherence to fundamental software and systems development and maintenance principles, most software development efforts have been significantly lacking at best. If the same inadequate practices are left unchanged in the face of increasing cyber threats, they will continue contributing to the insecure systems of tomorrow, just as they have contributed to the insecure systems of today.

The fundamental problem underlying this kind of assertion is the difficulty of justifying what are often perceived as non–revenue-generating activities and the need to get the product to market on or before delivery timelines. The difficulty also stems from the infeasibility of proving a negative. For example, one cannot prove to the uninvolved or unimpressed observer that the time, money, and human resources spent on prevention are well spent simply because a disaster or incident was averted. Such an observer will take the fact that it did not occur as evidence that the possibility never really existed, or that the problem was exaggerated. In the same way, a skeptical observer may take the absence of security breaches as evidence that no such threat existed, or that it was never as bad as claimed. The problem is exacerbated after spending time and effort on establishing security policies, security controls, and enforcement mechanisms.

In both cases, the money, time, and effort expended to prevent adverse computing consequences from harming an organization are viewed with suspicion and not qualifying at a level to justify the change in delivery schedule or funding. However, such suspicions are not new in the world of information processing. Throughout the history of IT, this attitude has existed regarding problem-prevention activities such as quality assurance, quality control testing, documentation, configuration management, change management, and other system-management controls designed to improve quality, reliability, and system maintainability.

If you consider the analogy of application development as being similar to the way a house is built, it is much harder to implement and integrate something after the fact. Security is not and should not be considered a superficial exercise similar to hanging the drapes or changing a light bulb. It is like the plumbing of a house, and it needs to be implemented with that in mind.

Problems with security are management problems. Management dictates how much time and money can be spent on security activities such as risk management and software testing. How much time and energy are invested in designing, building, testing, and fixing software is in the end a management call. Thus, the primary concern of those striving to improve software security must be the correction of these same basic and fundamental IT management defects; otherwise effective security can never be attained or maintained.

38.2.1 Software Total Quality Management. The nature of the computer security challenge is that of dynamically changing vulnerabilities and threats, based on rapidly changing technologies, used in increasingly complex business applications. Organizations need to build software around the idea of a continuous improvement model, as popularized by the total quality management (TQM) articulated in the ISO 9000 family of standards. For purposes of quality improvement, this model was viewed as a continuing cycle of plan-do-check-and-act. For software security, it can be thought of as a cycle of plan-fix-monitor-and-assess. A security management process to

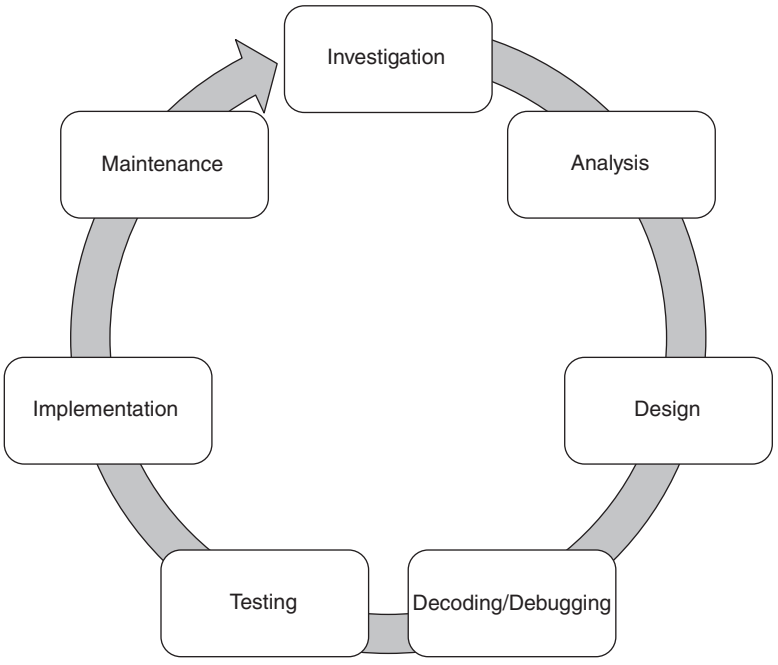


EXHIBIT 38.1 Common Lifecycle Process

execute this model must be established and integrated into the day-to-day operations of the business. Integration means that the security initiative does not stand alone as a watchdog or act merely as an audit function, and it must move away from a policing function. Security must be considered an essential part of all other systems development and management activities, including requirements definition, systems design, programming, package integration, testing, systems documentation, training, configuration management, operations, and maintenance. Security must be viewed as good architecture. In other words, security must be a part of the entire development lifecycle and management gates of the project. Exhibit 38.1 provides a common lifecycle process as described in Chapter 39.3.1.

A security management process must be designed in such a way that the activities of planning, fixing, monitoring, and assessing are accomplished in an iterative fashion for systems being defined, designed, and implemented and for systems that are currently executing the day-to-day business operations.

For many organizations, it has not been uncommon for software to be built with little or no attention to security concerns. Consequently, many business-critical applications have made their way into full production with inadequate or nonexistent security controls in place. For example, a quickly built application to resolve an initial need for tracking information within a database through a Web interface may not get all the access control configurations designed into the application that are needed for proper security. It serves a quick purpose but becomes a long-term tool because of its success. Because of the nature of business (moving from one task to the next), nothing is really done to develop it into the function it now fills. In addition, the longer the application stays in use, the greater the likelihood that the original developers are to transition, and any historical knowledge or documentation will be lost. This makes the ability to update or correct the security of the application less desirable. As a result, users or

hackers may be unintentionally or intentionally capable of accessing and modifying critical information within the database or within the system itself. Fortunately, this trend has started to change, but it is still in need of significant attention. Because such systems are already operating, the only way realistically to identify and assess their security risks outside of replacing the application outright with an application that includes the security requirements is to begin aggressively monitoring for intrusions and then to design and implement corrective security policies and controls based on what is discovered. The effective sequence of security activities for systems already in operational status will be to monitor current operations, assess suspicious activity, plan a corrective policy and technology control, implement the control on the system, and then monitor for continuing effectiveness.

38.2.2 Due Diligence. To address the software and code security effectively, it must become an issue for management, just as other security concerns have began to demand management attention. Business executives and IT managers are now all too familiar with the concept of due diligence as applied to the uses of information technology. Because of the potential for legal fallout, it has become necessary to view IT-related actions through the definitions of due diligence and reasonable care.

The significance of the concepts of due diligence and reasonable care is that they allow for an evolving metric against which an organization's application security deliberations, decisions, and actions can be compared. The comparison is usually against a similar organization, in like circumstances of vulnerability and threat, and with similar predictable adverse impacts on customers, partners, shareholders, employees, and the public.

For example, if one company employs a security control and does not experience any security breaches that the technique was supposed to prevent, that fact could be used to establish a baseline against which other similar organizations could be compared. If enough similar organizations employ the same technique, the security control may become categorized as a best practice for that industry.

If, however, another company in the same industry did not employ the security control and did experience security breaches of the type the technique was suppose to prevent, it might clearly indicate a lack of due diligence or reasonable care. Software security decisions and implementations are not a one-time event, but need to be under a continuous process of risk evaluation, management, and improvement. The decisions start with the selection of the development language and continue from there. The choice of an interpreted versus compiled language could have an impact not only to development skills required and timelines, but how the use of the languages libraries are use and the impact future updates can be accomplished by the customer. A good example is the recent[1] security issues that have plagued Java and Microsoft products.

It is therefore imperative, in order to demonstrate an ability to exercise continual due diligence to management, to establish a documented computer security risk management program, regardless of internal or external software development, and to integrate it into the overall management processes of the software development process. Nothing less will demonstrate that a company, and its board, is capable of assessing computer security threats and of acting in a reasonable manner.

For software developers working on commercial software for sale to a specific client under contract, or even to a wide range of customers governed by end-user license agreements, the exercise of due diligence and improved security may reduce

the risk of lawsuits claiming damages for negligence. Whether successful or not, such lawsuits are never positive publicity for makers of software products.

38.2.3 Regulatory and Compliance Considerations. Regulatory and compliance considerations have become a significant aspect to most businesses. The additional aspect of software development increases the need to be aware of the implication state, federal, or international regulation may have on a development project. Compliance with nongovernmental industries may also need to be considered.

If an organization is subject to specific regulations (e.g., Sarbanes-Oxley legislation), then documentation of the errors encountered, and the resulting internal reporting and remediation efforts, is critical. The error should clearly indicate how it was identified, who identified it, how it was reported to management, what the remediation will be, and when it is anticipated to be completed. Without these details, management may encounter significant difficulties in confirming that adequate internal control mechanisms exist, together with the appropriate and adequate involvement of management. In addition, an error could result in a control weakness being identified by the external auditors or, even worse, as a material weakness (a material weakness would be the use of a programming language that has a known inherent flaw rather than another more appropriate and without the flaw, similar to using aluminum instead of steel for a bridge), depending on its nature and severity.

Additionally, when errors are noted that affect multilocation systems or applications, the significance and materiality—that is to say the extent to which the errors impact or may impact the system or any other system to which the application interfaces—must be considered from the aspect of both the subsidiary and the headquarters locales.

Since laws differ among states and countries, what might be considered legally acceptable standards of privacy or accounting in one locale might not be acceptable elsewhere. As a result, careful consideration of the requirements and intended uses or customers must be established as part of the project initiation and due diligence processes.

38.3 TECHNICAL AND PROCEDURAL ISSUES. Security should be integrated into every stage of the application life cycle. This includes the requirements analysis, the design stage for software, and the operating system security kernel, in corporate policy development, and in human awareness, training, and education programs.

Writing secure code can be a daunting and technically challenging undertaking. That is not to say that it cannot be done, but it requires the developer to work in conjunction with the rest of the project team and other key stakeholders to meet the challenges that must be defined, reconciled, and overcome prior to the release of the software or system. The additional push to market that can drive the rapid development process can also have a negative impact on the security of the application. The problem must be dealt with in two ways, technical and procedural. Technically, the programmers must be aware of the pitfalls associated with application development and avoid them. Procedurally, the development team and organization need to adhere to a consistent methodology of development. This consistency also needs to include the identification of security risks at every stage of the process, including the requirements analysis.

38.3.1 Requirements Analysis. Security must be a part of the requirements analysis within any development project. Most application developers know that adding

security after the fact increases cost and time, and becomes more complicated. Requirements can be derived from different sources:

- Functional needs of the system (this includes the operating system the application will run on)
- National, international, or organizational standards or guidelines
- Regulatory restrictions
- Sensitivity level of data
- Existing security policies
- Cost/benefit analysis
- Risk rating

The purpose of analysis is to determine what information and processes are needed to support the desired objectives and functions of the software and the business model.

38.3.2 Design. The informational data gathered during the analysis gathering goes into the software design as requirements. What comes out of the analysis are the data, logic, and procedural design.

Developers take the data and the informational model data and transform them into the data structures that will be required to implement the software. The logic design defines the relationships between the major structures and components of the application. The procedural design transforms structural components into descriptive procedures. Access control mechanisms are chosen, rights and permissions are defined, any other security specifications are appraised, and solutions are determined. A work breakdown structure includes the development and implementation stages. The structure includes a timeline and detailed activities for testing, development, staging, integration testing, and product delivery.

The decisions made during the design phase are pivotal to application development. The design is the only way the requirements are translated into software components. It is in this way that software design is the foundation of the development process, and it greatly affects software quality and maintenance. If good product design is not put in place in the beginning, the rest of the development process will be that much more challenging.

38.3.3 Operating System. The operating system security kernel is responsible for enforcing the security policy within the operating system and the application. As such, the architecture of the kernel operating system is typically layered, with the kernel at the most privileged layer. This is a small portion of the operating system, but all references to information and changes to authorization pass through the kernel.

To be secure the kernel must meet three basic criteria:

1. **Completeness.** All access to information must go through the kernel.
2. **Isolation.** The kernel must be protected from unauthorized access.
3. **Verifiability.** The kernel must be proven to meet design specifications.

As mentioned, the kernel runs in the most privileged layer. Most operating systems have two processor access modes, user and kernel. General application code runs in user mode, while the operating system runs in kernel mode. Kernel mode allows

the processor full access to all system memory and CPU instructions. When applications are not written securely, this separation of modes can become compromised, enabling exploitation of vulnerabilities through arbitrary code, buffer overflows, and other techniques.

38.3.4 Best Practices and Guidelines. "Best practices" is a term that can cause vigorous debate, especially with regard to security. Best practices in general, and particularly with regard to security, often fall prey to a dogmatic, and sometimes blind, devotion to nonsensical practices that have little to do with security and more to do with faith or tradition. Regardless, best practices help provide a set of guidelines that help provide structure. In addition, best practices can be adapted to help meet the needs of a particular situation. An example would be the National Institute of Standards and Technology (NIST) Special Publication Series 800. This series provides best practices and recommendations that can be adapted or integrated into other practices to assist in improving and eliminating the tradition to a particular practice that may no longer have a legitimate use. Considering this, most general security textbooks contain recommendations on security-related aspects of programming—see, for example, Stallings—that do in fact have a very real benefit in creating more secure software.

In addition to designing security into a system from the start, there are also some obvious guidelines that will help in developing secure software:

- Impose strong identification and authentication (I&A) for critical and sensitive systems in addition to the I&A available from the operating system; ideally, use token-based or biometric authentication as part of the initialization phase of your application.
- Document your code thoroughly, including using data dictionaries for full definition of allowable input and output to functions and allowable range and type of values for all variables.
- Use local variables, not global variables, when storing sensitive data that should be used only within a specific routine (i.e., use the architecture of the process stack to limit inadvertent or unauthorized access to data in the stack).
- Reinitialize temporary storage immediately after the last legitimate use for the variable, thus making scavenging harder for malefactors.
- Limit functionality in a specific module to what is required for a specific job (e.g., do not use the same module for supervisory functions and for routine functions carried out by clerical staff).
- Define views of data in databases that conform to functional requirements and limit access to sensitive data (e.g., the view of data from a medical records database should exclude patient identifiers when a worker in the finance department is using the database for statistical aggregation).
- Use strong encryption (*not* homegrown encryption) that has industry-standard routines to safeguard sensitive and critical data on disk. Locally developed, homegrown encryption is generally not as safe.
- Disallow access by programmers to production databases.
- Randomize or otherwise mask sensitive data when generating test subsets from production data.

- Use test-coverage monitors to verify that all sections of source code are in fact exercised during quality assurance tests; investigate the functions of code that never is executed.
- Integrate logging capability into all applications for debugging work, for data recovery after crashes in the middle of a transaction, and for security purposes such as forensic analysis.
- Create log-file records that include a cryptographically sound message authentication code (MAC) that itself includes the MAC of the preceding record as input for the algorithm; this technique ensures that forging a log file or modifying it will be more difficult for a malefactor.
- Log all process initiations for a program and log process termination; include full details of who loaded the program or module.
- Log all modifications to records and optionally provide logging for read access as well.
- Use record-level locking to prevent inadvertent overwriting of data on records that are accessed concurrently. Be sure to unlock a sequence of locks in the inverse order of the lock sequence to prevent deadlocks. (Thus, if you lock resource A, B, and C in that order, unlock C, then B, then A.)
- Sign your source code using digital signatures.
- Use checksums in production executables to make unauthorized modifications more difficult to conceal.

Mike Gerdes, former manager at AtomicTangerine, contributed these suggestions in the course of a discussion:

- Recommend the readers adopt a practice of designing code in a more holistic fashion. A common practice is to write and test routines in a way that verifies the code processes the data in the way intended. To avoid the effects of malicious code and data input attacks, the programmer must also write code that deals with what is *not* supposed to be processed. A more complete design methodology would also include testing of all inbound information to ensure exclusion of any data which did not fit the requirements for acceptable data. This method should be applied to high-risk applications and those with an extremely arduous test cycle and will eliminate many of the common attack methods used today.
- Establish the criteria for determining the sensitivity level of information contained in, or processed by, the application and subroutines.
- If they are not already present, consider implementing formal control procedures in the software programming methodology to ensure that all data is reviewed during QA processes and to be sure it is classified and handled appropriately for the level assigned.
- Identify and include any mandatory operating system, and network security characteristics for the production system in the specifications of the software. In addition to providing the development and QA teams some definition of the environment the software is designed to run in, giving the administrator and end users an idea of what your expectations were when you created the code can be extremely useful in determining where software can, or cannot, be used.

- Where appropriate, verify the digital signatures of routines that process sensitive data when the code is being loaded for execution.
- If you include checksums on executables for production code, include routines that verify the checksums at every system restart.

The Carnegie-Mellon Software Engineering Institute Computer Emergency Response Team (CERT) provides a good "Top 10 Secure Coding Practices"[2] list. The list below provides some additional items to consider in addition to the items previously listed:

- Validate input from data sources. Proper input validation can eliminate many software vulnerabilities.
- Compile code using the highest warning level available for the compiler and eliminate warnings by modifying the code.
- Create a software architecture and design the software to implement and enforce security. Carnegie-Mellon has pre-established programming standard by language that can be found at www.cert.org/secure-coding/scstandards.html
- Remember the K-I-S-S (Keep It Simple, Stupid) principle. Complex designs increase the likelihood of errors or vulnerabilities and the need for security mechanisms or resulting mitigations to become more complex.
- Base access decisions on permission rather than exclusion to adhere to the principle of least privilege.
- Sanitize data sent to other systems, such as command shells, relational databases, and commercial off-the-shelf (COTS) components.
- Practice defense in depth within the application, so that if one layer of defense turns out to be inadequate, another layer of defense can prevent or limit a security flaw from becoming an exploitable vulnerability.
- Establish good quality assurance techniques to be more effective in identifying and eliminating vulnerabilities.
- Develop and/or apply a secure coding standard for your target development language and platform.
- Define security requirements at the start of the project.
- Think like a hacker and use threat modeling to anticipate the threats to which the software may be subjected. If you consider how it can be attacked, you will be able to prevent the attack before it can occur.

38.3.5 Languages. To date, no common computer languages have security specific features built-in. Java does include provisions for limiting access to resources outside the "sandbox" reserved for a process, as described in the books by Felten and McGraw. Nevertheless, Java is still one of the languages targeted as an attack vector by many malware implementations. PASCAL uses strong typing and requires full definition of data structures, thus making it harder to access data and code outside the virtual machine defined for a given process. In contrast, C and C++ allow programmers to access any region of memory at any time the operating system permits it.

Computer languages allow developers to write code as well as they can. Strongly typed languages may offer better constraints on programmers, but the essential

requirement is that the programmers continue to think about security as they design and build code.

There are several sets of security utilities and resources available for programmers; for example, RSA has a number of cryptographic toolkits. Some textbooks (e.g., Schneier's *Applied Cryptography*) include CD-ROMs with sample code. In addition, the Computer Emergency Response Team Coordination Center (CERT/CC) was started in December 1988 by the Defense Advanced Research Projects Agency (DARPA), which was part of the U.S. Department of Defense. CERT/CC is located at the Software Engineering Institute, a federally funded research center operated by Carnegie Mellon University.

CERT/CC studies Internet security vulnerabilities, provides services to Websites that have been attacked, and publishes security alerts. CERT/CC's research activities include the area of WAN computing and developing improved Internet security.

As mentioned in Section 38.3.4, Carnegie-Mellon has established secure coding standards based on language that can help establish a uniform set of principles and guidelines across your development group. Regardless of the language selected for software development, it is important to make the decision based on the strengths and weaknesses of the language. This means that a risk analysis needs to be a part of the design of any development project, including the selection of the development language. By understanding those strengths and weaknesses, better decisions can be made to prevent the inherent vulnerabilities within the language from becoming exposed inadvertently during development.

38.4 TYPES OF SOFTWARE ERRORS.

New programmers should review the range of root causes for software errors. Such review is particularly useful for students who have completed training that did not include discussions of systematic quality assurance methodology. The next sections can serve as a basis for creating effective sets of test data and test procedures for unit tests of new or modified routines.

38.4.1 Internal Design or Implementation Errors.

A general definition of a software error is a mismatch between a program and its specifications; a more specific definition is the failure of a program to do what the end user reasonably expects. There are many types of software errors. Some of the most important include:

- Initialization
- Logic flow
- Calculation
- Boundary condition violations
- Parameter passing
- Race condition
- Load condition
- Resource exhaustion
- Resource, address, or program conflict with the operating system or application(s)
- Regulatory compliance considerations
- Other errors

38.4.1.1 Initialization. Initialization errors are insidious and difficult to find. The most insidious programs save initialization information to disk and fail only the first time used—that is, before they create the initialization file. The second time a given user activates the program, there are no further initialization errors. Thus, the bugs appear only to employees and customers when they activate a fresh copy of the defective program. Other programs with initialization errors may show odd calculations or other flaws the first time they are used or initialized; because they do not store their initialization values, these initialization errors will continue to reappear each time the program is used.

38.4.1.2 Logic Flow. Modules pass control to each other or to other programs. If execution passes to the wrong module, a logic-flow error has occurred. Examples include calling the wrong function, or branching to a subroutine that lacks a RETURN instruction, so that execution falls through the logical end of a module and begin executing some other code module.

38.4.1.3 Calculation. When a program misinterprets complicated formulas and loses precision as it calculates, it is likely that a calculation error has occurred; for example, an intermediate value may be stored in an array with 16 bits of precision when it needs 32 bits. This category of errors also includes computational errors due to incorrect algorithms.

38.4.1.4 Boundary Condition Violations. Boundaries refer to the largest and smallest values with which a program can cope; for example, an array may be dimensioned with 365 values to account for days of the year, and then fail in a leap year when the program increments the day-counter to 366 and thereby attempts to store a value in an illegal address. Programs that set variable ranges and memory allocation may work correctly within the boundaries but, if incorrectly designed, may crash at or outside the boundaries. The first use of a program also can be considered a boundary condition.

One of the most important types of boundary violations is the buffer overflow. In this error, data placed into storage exceed the defined maximum size and overflow into a section of memory identified as belonging to one or more different variables. The consequences can include data corruption (e.g., if the overflow overwrites data that are interpreted as numerical or literal values) or changes in the flow of execution (e.g., if the altered data include logical flags that are tested in branch instructions).

Buffer overflows have been exploited by writers of malicious code who insert data that overflows into memory areas of interpreted programs and thus become executed as code.

All programs should verify that the data being stored in an array or in any memory location do not exceed the expected size of the input. Data exceeding the expected size should be rejected or, at least, truncated to prevent buffer overflow.

38.4.1.5 Parameter Passing. Sometimes there are errors in passing data back and forth among modules. For instance, a call to a function accidentally might pass the wrong variable name so that the function acts on the wrong values. When these parameter-passing errors occur, data may be corrupted, and the execution path may be affected because of incorrect results of calculations or comparisons. As a result, the

latest changes to the data might be lost, or execution might fall into error-handling routines even though the intended data were correct.

38.4.1.6 Race Condition. When a race occurs between event A and event B, a specific sequence of events is required for correct operation, but the program does not ensure this sequence. For example, if process A locks resource 1 and waits for resource 2 to be unlocked while process B locks resource 2 and waits for resource 1 to be unlocked, there may be a deadly embrace that freezes the operations if the processes overlap in execution. If they do not happen to overlap, there is no problem at that time.

Race conditions can be expected in multiprocessing systems and interactive systems, but they can be difficult to replicate; for example, the deadly embrace just described might happen only once in 1,000 transactions if the average transaction time is very short. Consequently, race conditions are among the most difficult to detect during quality assurance testing and are best identified in code reviews. Programmers should establish and comply with standards on sequential operations that require exclusive access to more than one resource. For example, if all processes exclusive-lock resources in a given sequence and unlock them in the reverse order, there can be no deadly embrace.

38.4.1.7 Load Condition. All programs and systems have limits to storage capacity, numbers of users, transactions, and throughput. Load errors are caused by exceeding the volume limitations of storage, transactions, users, and networks can occur due to high volume, which includes a great deal of work over a long period, or high stress, which includes the maximum load all at one time. For example, if the total theoretical number of transactions causes a demand on the disk I/O system that exceeds the throughput of the disk controller, processes will necessarily begin piling up in a queue waiting for completion of disk I/Os. Although theoretical calculations can help to identify where possible bottlenecks can occur in CPU, memory, disk, and network resources, a useful adjunct is automated testing that permits simulation of maximum loads defined by service-level agreements.

38.4.1.8 Resource Exhaustion. The program running out of high-speed memory (RAM), mass storage (disk), central processing unit (CPU) cycles, operating system table entries, semaphores, network bandwidth, or other resources can cause failure of the program. For example, inadequate main memory may cause excessive swapping of data to disk (thrashing), typically causing drastic reductions in throughput, because disk I/O is typically 1,000 times slower than memory access.

38.4.1.9 Interapplication Conflicts. With operating systems (OS) as complex as they are, OS manufacturers routinely distribute the code requirements and certain parameters to the application software manufacturers, so that the likelihood of program conflicts or unexpected stoppages are minimized. While this certainly helps reduce the number of problems and improves the forward and backward compatibility with previous OS versions, even the OS vendors on occasion experience or cause difficulties when they do not conform to the parameters established for their own programs.

38.4.1.10 Other Sources of Error. It is not unusual for errors to occur where programs send bad data to devices, ignore error codes coming back, and even try to use devices that are busy or missing. The hardware might well be broken, but the

software also is considered to be in error when it does not recover from such hardware conditions.

Additional errors can occur through improper builds of the executable; for example, if an old version of a module is linked to the latest version of the rest of the program, the wrong sign-on screens may pop up, the wrong copyright messages may be displayed, the wrong version numbers may appear, and various other inaccuracies may occur.

38.4.1.11 User Interface. Generally speaking, the term "user interface" denotes all aspects of a system that are relevant to a user. It can be broadly described as the user virtual machine (UVM). This would include all screens, the mouse and keyboard, printed outputs, and all other elements with which the user interacts. A major problem arises when system designers cannot put themselves in the user's place and cannot foresee the problems that technologically challenged users will have with an interface designed by a technologically knowledgeable person.

Documentation is a crucial part of every system. Each phase of development—requirements, analysis, development, coding, testing, errors, error solutions and modifications, implementation, and maintenance—needs to be documented. All documents and their various versions need to be retained for both future reference and auditing purposes. Additionally, it is important to document the correct use of the system and to provide adequate instructional and reference materials to the user. Security policies and related enforcement and penalties also need to be documented. Ideally, the documentation should enable any technically qualified person to repair or modify any element, as long as the system remains operational.

38.4.1.12 Functionality. A program has a functionality error if performance that can reasonably be expected is confusing, awkward, difficult, or impossible. Functionality errors often involve key features or functions that have never been implemented. Additional functionality errors exist when:

- Features are not documented.
- Required information is missing.
- A program fails to acknowledge legitimate input.
- There are factual errors or conflicting names for features.
- There is information overload.
- The material is written to an inappropriate reading level.
- The cursor disappears, or is in the wrong place.
- Screen displays are wrong.
- Instructions are obscured.
- Identical functions require different operations in different screens.
- Improperly formatted input screens exist.
- Passwords or other confidential information are not obscured or protected adequately.
- Tracing the user data entry or changes is unavailable or incomplete.
- Segregation of duties is not enforced. (This can be particularly critical for organizations subject to legal and regulatory requirements.)

38.4.1.13 Control (Command) Structure. Control structure errors can cause serious problems because they can result in:

- Users getting lost in a program.
- Users wasting time because they must deal with confusing commands.
- Loss of data or the unwanted exposure of data.
- Work delay.
- Financial cost.
- Unanticipated exposure to data leakage or compromise; this can result in significant liability if consumers' personal identifying information (PII) is compromised.
- Data not being encrypted as intended or being visible to unauthorized users.

Some common errors include:

- Inability to move between menus.
- Confusing and repetitive menus.
- Failure to allow adequate command-line entries.
- Requiring command-line entries that are neither intuitive nor clearly defined on screen.
- Failure of the application program to follow the operating system's conventions.
- Failure to distinguish between source and parameter files, resulting in the wrong values being made available to the user through the interface, or failure to identify the source of the error
- Inappropriate use of the keyboard, when new programs do not meet the standard of a keyboard that has labeled function keys tied to standard meanings.
- Missing commands from the code and screens resulting in the user being unable to access information, to utilize programs, or to provide for the system to be backed up and recoverable. There are a host of other commands that can leave the system in a state of less-than-optimum operability.
- Inadequate privacy or security that can result in confidential information being divulged, in the complete change or loss of data without recoverability, in poor reporting, and even in undesired access by outside parties.

38.4.1.14 Performance. Speed is important in interactive software. If a user feels that the program is working slowly, that can be an immediate problem.

Performance problems include slow response, unannounced case sensitivity, uncontrollable and excessively frequent automatic saves, inability to save, and limited scrolling speed.

There are five fundamental potential bottlenecks to consider in analyzing any computer and network performance problem relating to access to and speed of:

- The central processing unit(s)
- Main memory
- Secondary memory (e.g., disk drives)

- Application design
- Network bandwidth

Slow operation can depend on (but is not limited to) the OS, the other applications running, memory saturation and thrashing (excessive swapping of memory contents to virtual memory on disk or on other, relatively slow, memory resources such as flash drives), memory leakage (the failure to de-allocate memory that is no longer needed), disk I/O inefficiencies (e.g., reading single records from very large blocks), and program conflicts (e.g., locking errors). For more information on program performance.

At another level, performance suffers when program designs make it difficult to change their functionality in response to changing requirements. In a database design, defining a primary index field that determines the sequence in which records are stored on disk can greatly speed access to records during sequential reads on key values for that index—but it can be counterproductive if the predominant method for accessing the records is sequential reads on a completely different index.

38.4.1.15 Output Format. Output format errors can be frustrating and time consuming. An error is considered to have occurred when the user cannot change fonts, underlining, boldface, and spacing that influence the final look of the output; alternatively, delays or errors when printing or saving document may occur. Errors occur when the user cannot control the content, scaling, and look of tables, figures, and graphs. Additionally, there are output errors that involve expression of the data to an inappropriate level of precision.

38.5 ASSURANCE TOOLS AND TECHNIQUES.

The security community has known about buffer overflow vulnerabilities, as an example, for the past 25 or 30 years, yet the vulnerabilities are still increasing. For example, the National Vulnerability Database reported 418 new buffer overflow errors in 2013 compared with 392 in 2011. Clearly, greater efforts are necessary to teach and to enforce awareness, implementation, and testing of source code.

38.5.1 Education Resources.

Build Security In (BSI) is a project of the Software Assurance program of the Strategic Initiatives Branch of the National Cyber Security Division (NCSD) of the U.S. Department of Homeland Security. BSI (http://buildsecurityin.us-cert.gov) contains and links to best practices, tools, guidelines, rules, principles, and other resources that software developers, architects, and security practitioners can use to build security into software in every phase of its development. BSI content is based on the principle that software security is fundamentally a software engineering problem and must be addressed in a systematic way throughout the software development life cycle.

In addition to BSI, and out of the multitudes of sites on the Web that address some facet of secure coding, some sites are listed next to assist in secure application development:

- Carnegie-Mellon Software Engineering Institute: www.cert.org/secure-coding
- FreeBSD Security Information: www.freebsd.org/security/security.html
- International Systems Security Engineering Association (ISSEA): www.issea.org
- Open Web Application Security Project: www.owasp.org

- Secure, Efficient, and Easy C Programming: http://irccrew.org/~cras/security/c-guide.html
- Secure Programming for Linux and UNIX HOWTO: http://www.dwheeler.com/secure-programs/
- Windows Security: www.windowsecurity.com
- World Wide Web Security FAQ: http://www.w3.org/Security/Faq/

38.5.2 Code Examination and Application Testing Techniques. Scanning once is not enough; ongoing application assessment is essential to implementing effective secure development practices and, in turn, a secure application. Like many other aspects of information technology and computer security, a layered approach to testing is also needed. As a result, there are numerous techniques and approaches, both manual and automated, to software security testing. This section will provide a cursory overview of the major techniques and approaches.

38.5.2.1 Static Analysis. Static analysis or source code review is the process of manually checking source code for security weaknesses. Sometimes looking at the code is the only way to resolve issues. Understanding the intent of the inner workings of the application can identify serious security issues that cannot be identified any other way.

Some issues that are particularly conducive for discovery through static analysis include concurrency problems, flawed business logic, access control problems, backdoors, Trojans, Easter eggs, time/logic bombs, and other forms of malicious code.

38.5.2.2 Binary Code Analysis. Binary code analysis utilizes tools that support analysis of binary executables such as decompilers, disassemblers, and binary code scanners. This is to say that binary analysis is the reverse engineering of the executable. The least intrusive technique is binary scanning. Binary scanners analyze machine code to model a language-neutral representation of the program's behaviors, control/data flows, call trees, and external function calls. An automated vulnerability scanner then traverses the model in order to locate vulnerabilities caused by common coding errors and simple back doors.

The next least intrusive technique is disassembly of the binary. The binary code is reverse engineered to an intermediate assembly language, but the disadvantage of disassembly is that it requires an expert who both thoroughly understands that particular assembler language and who is skilled in detecting security-relevant constructs within assembler code.

The most intrusive reverse engineering technique is decompilation of the code. The binary code is reverse engineered completely to source code, which can then be subjected to the same security code review techniques and other white box tests as original source code. Note, however, that decompilation is technically problematic in that the quality of the source code generated through decompilation is often very poor. As a result, code is rarely as navigable or comprehensible as the original source code. The analysis of decompiled source code will always be more difficult and time consuming than that of the original source code. For this reason, the use of decompilation for security analysis should only be used for the most critical systems.

Another point to consider regarding reverse engineering is that it may also be legally prohibited. A majority of software vendors' license agreements prohibit reverse engineering. In addition, many software vendors repeatedly cite the Digital Millennium

Copyright Act (DMCA) of 1999 as justification, even though the DMCA explicitly exempts reverse engineering and encryption research from the prohibitions against copy protection circumvention.

38.5.2.3 White Box Approach. White box testing strategy deals with the internal logic and structure of the code. White box testing is also known as glass, structural, open box, or clear box testing. In order to implement white box testing, the tester has to deal with the code, and therefore needs to possess knowledge of coding and logic (i.e., internal working of the code). White box tests also need the tester to look into the code and to find out which unit/statement/chunk of the code is malfunctioning.

Advantages of white box testing are:

- It becomes very easy to find out which type of input/data can help in testing the application effectively.
- It helps in optimizing the code.
- It helps in removing the extra lines of code, which can bring in hidden defects.

Disadvantages of white box testing are:

- A skilled tester is needed to carry out this type of testing, which increases the cost.
- It is nearly impossible to look into every bit of code to find hidden errors that may create problems, resulting in failure of the application.

38.5.2.4 Black Box Approach. Black box testing is testing without knowledge of the internal workings of the item being tested. Black box testing is also known as behavioral, functional, opaque box, and closed box. For this reason, the tester and the programmer can be independent of one another, avoiding programmer bias toward his own work. Test groups are often used. Due to the nature of black box testing, the test planning can begin as soon as the specifications are written.

Advantages of black box testing are:

- It is more effective on larger units of code than glass box testing.
- The tester needs no knowledge of implementation, including specific programming languages.
- Tester and programmer are independent of each other.
- Tests are done from a user's point of view.
- It will help to expose any ambiguities or inconsistencies in the specifications.
- Test cases can be designed as soon as the specifications are complete.

Disadvantages of black box testing are:

- Only a small number of possible inputs can actually be tested; to test every possible input stream on a complex system would take nearly forever.
- Without clear and concise specifications, test cases are hard to design.
- There may be unnecessary repetition of test inputs if the tester is not informed of test cases the programmer has already tried. It is always critical to validate any

testing claims through reports or result output to help avoid duplication. Likewise, it is always important to rerun tests that may have ambiguous results.

- It may leave many program paths untested.
- It cannot be directed toward specific segments of code that may be very complex.

38.5.2.5 Gray Box—Blended Approach. Gray box testing is a software testing technique that uses a combination of black box testing and white box testing. Gray box testing is not black box testing, because the tester does know some of the internal workings of the software under test. In gray box testing, the tester applies a limited number of test cases to the internal workings of the software under test. In the remaining part of the gray box testing, the tester takes a black box approach in applying inputs to the software under test and observing the outputs.

38.5.2.6 Fault Injection. There are two primary categories of fault injection: source code or binary fault injection. Fault injection is used to induce stress in the software, create interoperability problems among components, simulate faults in the execution environment, and thereby revealing safety-threatening faults that are not made apparent by traditional testing techniques. Security extends standard fault injection by adding error injection, enabling testers to analyze the security impacts of the behaviors, and state changes resulting in the software when exposed to changes in the environment data. These data changes are meant to simulate the types of faults that would occur during unintentional user errors or intentional attacks on the software through its environment, including attacks on the environment itself. A data change is simply an alteration of the data the execution environment passes to the software, or between software components. Fault injection reveals both the effects of security faults on individual component behaviors, and the behavior of the system as a whole.

Binary fault injection is useful when performed as part of penetration testing to enable the tester to obtain a more holistic picture of how the software responds to attacks. However, faults should not be limited to simulating real world attacks. Fault injection scenarios should be designed to give a complete understanding of the security impacts that the application's behaviors, states, and properties exhibit under all possible operating conditions.

The main challenges in fault injection testing are determining the most meaningful combination and volume of faults to inject and, as with all tests, interpreting the results. Furthermore, just because fault injection does not cause the software to behave in a nonsecure manner or fail in a nonsecure state, this cannot be interpreted to mean that the software will perform equally when exposed to the more complex inputs it typically receives during execution.

38.5.2.7 Fuzzing. Fuzz testing or fuzzing is a technique that is often automated or semi-automated, that involves providing invalid, unexpected, or random data to the inputs of a computer program. The program is then monitored for exceptions, failing built-in code assertions, or for finding potential memory leaks. There are two forms of fuzzing, mutation-based and generation-based, which can be employed as part of white, black, or grey box testing.

The idea behind fuzz testing is to look for interesting program behavior that results from noise injection and may indicate the presence of vulnerabilities or faults. Fuzzing

applications, fuzzers, are generally specific to a particular type of input and written to test a specific program; they are not easily reusable, but file formats and network protocols are the most common targets of testing. Any type of program input can be fuzzed, including items not normally considered "input" (e.g., databases, shared memory, etc.). The value is the specificity, that is to say, fuzzing can often reveal security vulnerabilities that generic testing tools such as vulnerability scanners and fault injectors cannot. On the other hand, completely random fuzzing is a typically ineffective way to uncover problems in an application, and as a result fuzzing technology has evolved to include more intelligent techniques.

38.5.2.8 Vulnerability Scanning. Automated vulnerability scanning is an automated scanning technique for application level software, as well as for Web servers, database management systems, and some operating systems. The tools scan the application software for input and output of known patterns associated with known vulnerabilities. The vulnerability patterns, or signatures, are similar to the signatures used by virus scanners, or the coding constructs searched for by automated source code scanners, making the vulnerability scanner, essentially, an automated pattern-matching tool. Some Web application vulnerability scanners additionally attempt to perform automated stateful assessment using simulated reconnaissance attack patterns and fuzzing techniques to probe for known or common vulnerabilities.

Like any signature-based scanner, vulnerability scanners can report false positives. As a result, the tester must have enough software and security expertise to meaningfully interpret the scanner's results to limit the false positives or negatives, to prevent the creation of an incident where it does not exist. This is why it is important to use a layered approach and combine different test techniques to examine the software for vulnerabilities.

Vulnerabilities in software are often masked by environmental protections such as network and application-level firewalls. Furthermore, environment conditions may create unique vulnerabilities that cannot be found by a signature-based tool, but using a combination of other black box tests, most especially penetration testing of the deployed software in the actual target environment. Vulnerability scanners are most effective as part of:

- Security assessments of binary components;
- Or before penetration testing (to eliminate the need to run penetration test scenarios for common vulnerabilities).

38.5.2.9 Software Penetration Testing. Penetration testing has been a common technique to test network security for many years and is also known as black box testing and ethical hacking. As described in the black box approach, penetration testing is essentially the testing of a running application (without knowing the inner workings of the application) to identify vulnerabilities. In penetration testing, the whole software system is a "live" environment. Typically, the penetration team would have access to the application like typical users. The testers then act like an attacker, attempting to find and exploit potential vulnerabilities.

Penetration testing should focus on aspects of system behavior, interaction, and vulnerability that cannot be observed through other testing techniques performed outside of a live environment. That means that the penetration should subject the system to

sophisticated multi-pattern attacks designed to trigger complex behaviors across system components. In addition, the testing should attempt to find security problems likely to originate in the architecture and design, rather than coding flaws, since this type of problem tends to be overlooked by other testing techniques. This also means test plan should include worst-case scenarios that reproduce threat vectors considered highly damaging, such as an insider threat scenario.

Many people use application penetration testing as the primary testing technique and while it has a place in a testing program; it should not be considered the primary or only testing technique. Penetration testing can potentially lead to a false sense of security. Just because an application passes a penetration test, it does not mean it is vulnerability free. Likewise, if an application fails, it is a strong indication that there are serious problems with it. However, focused penetration testing, which is testing that attempts to exploit vulnerabilities detected in previous tests, can be useful in detecting if those vulnerabilities are actually fixed within the code.

38.5.3 Standards and Best Practices. Testing is a complete software engineering discipline in its own right. Volumes have been written about the various techniques that software engineers use to test and validate their applications. Some basic best practices can be employed regardless of the testing methodology selected. These are:

- Perform automated testing.
- Make a test plan.
- Follow a specific methodology.
- Test at every stage.
- Test all system components.

In addition to these best practices, use of these standards can provide additional resources:

- ISO 17799, Information Technology: Code of Practice for Information Security Management
- ISO/IEC 15408, Evaluation Criteria for IT Security (the Common Criteria)
- SSE-CMM, System Security Engineering Capability Maturity Model
- ISO/IEC WD 15443, Information Technology: Security Techniques

38.6 CONCLUDING REMARKS. Computer security practitioners and top-level managers must understand that IT management deficiencies, left unchanged, will sabotage efforts to establish and sustain effective computer security programs. Furthermore, it is incumbent on those coding to help steer the project and educate management toward a model that is inclusive of security throughout the development life cycle.

The costs of neglecting security at the start will continue to be a serious issue until software security is fully integrated into every aspect of the development life cycle. Security must be built in, and building security into systems starts with the software development teams, practices, and techniques used to build those systems.

38.7 FURTHER READING

Campanella, J., ed. *Principles of Quality Costs: Implementation and Use*, 3rd ed. Milwaukee, Wisconsin: ASQ Quality Press, 1999.

Eilam, E. *Reversing: Secrets of Reverse Engineering*. Indianapolis, IN: Wiley Publishing, 2005.

Felten, E., & G. McGraw. *Securing Java: Getting Down to Business with Mobile Code*. New York: John Wiley & Sons, 1999. Also free and unlimited Web access from www.securingjava.com

Fox, C., and P. Zonneveld. *IT Control Objectives for Sarbanes-Oxley: The Role of IT in the Design and Implementation of Internal Control Over Financial Reporting*, 2nd ed. IT Governance Institute, 2006.

Hoglund, G., and G. McGraw. *Exploiting Software: How to Break Code*. Boston: Addison-Wesley, 2004.

Institute of Electrical and Electronics Engineers. *IEEE Standard for Software Test Documentation*. ANSI/IEEE Std. 829-1983. 1983.

Institute of Electrical and Electronics Engineers. *IEEE Standard for Software Quality Assurance Plans*. ANSI/IEEE Std. 730-1981. 1981.

Long, F., D. Mohindra, R. C. Seacord, D. F. Sutherland, and D. Svoboda. *Java Coding Guidelines: 75 Recommendations for Reliable and Secure Programs*. Addison-Wesley Professional, 2013

Martin, R. *et al. Clean Code: A Handbook of Agile Software Craftmanship*. New York: Prentice-Hall, 2008.

McConnell, S. *Code Complete*, 2nd ed. Redmond, WA: Microsoft Press, 2004.

McGraw, G., and E. W. Felten. *Java Security: Hostile Applets, Holes and Antidotes—What Every Netscape and Internet Explorer User Needs to Know*. New York: John Wiley & Sons, 1997.

McGraw, G., and E. W. Felten. " Understanding the Keys to Java Security—the Sandbox and Authentication," *JavaWorld*. May 1, 1997. www.javaworld .com/javaworld/jw-05-1997/jw-05-security.html.

Markantonakis, K., and K. Mayes, eds. *Secure Smart Embedded Devices, Platforms and Applications*. Springer, 2013.

NASA Software Assurance Technology Center. http://satc.gsfc.nasa.gov (URL inactive)

Ould, M.A. *Strategies for Software Engineering: The Management of Risk and Quality*. New York: John Wiley & Sons, 1990.

Open Web Application Security Project (OWASP), www.owasp.org

RSA Data Security, www.rsasecurity.com/products

Sebesta, R. *Concepts of Programming Languages*. 9th ed. Boston: Pearson Education, 2010.

Schneier, B. *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, 2nd ed. New York: John Wiley & Sons, 1995.

Secord, R. C. *Secure Coding in C and C++*, 2nd ed. Addison-Wesley Professional, 2013

Shore, J., and J. Warden. *The Art of Agile Development*. Sebastopol: O'Reilly Media, 2008.

Stallings, W. *Network and Internetwork Security: Principles and Practice*. Englewood Cliffs, NJ: Prentice Hall, 1995.

Tipton, H. F., and K. Henry. *Official (ISC) Guide to the CISSP CBK*. Boca Raton, FL: Auerbach Publications, 2006.

Wilhelm, T. *et al. Ninja Hacking: Unconventional Penetration Testing Tactics and Techniques*. New York: Syngress, 2011.

38.8 NOTES

1. Consider the following articles: "Ms. Smith," "Microsoft Mega-Patch Closes Critical IE Flaws, Fixes 57 Vulnerabilities," NetworkWorld, February 12, 2013, www.networkworld.com/community/blog/microsoft-mega-patch-closes-critical-ie-hole-fixes-57-vulnerabilities?source=NWWNLE_nlt_security_2013-02-13 or Danielle Walker, "Oracle Speaks, Promises To Get Java "Fixed Up." SC Magazine, January 28, 2013, www.scmagazine.com/oracle-speaks-promises-to-get-java-fixed-up/article/277898/?DCMP=EMC-SCUS_Newswire
2. CERT, "Top 10 Secure Coding Practices," last edited by Robert Seacord on March 1, 2011, www.securecoding.cert.org/confluence/display/seccode/Top+10+Secure+Coding+Practices

CHAPTER 39

SOFTWARE DEVELOPMENT AND QUALITY ASSURANCE

Diane E. Levine, John Mason, and Jennifer Hadley

39.1 INTRODUCTION 39·2

39.2 GOALS OF SOFTWARE QUALITY ASSURANCE 39·2

39.2.1 Uncover All of a Program's Problems 39·2
39.2.2 Reduce the Likelihood that Defective Programs Will Enter Production 39·3
39.2.3 Safeguard the Interests of Users 39·3
39.2.4 Safeguard the Interests of Software Producers 39·3

39.3 SOFTWARE DEVELOPMENT LIFE CYCLE 39·3

39.3.1 Phases of the Traditional Software Development Life Cycle 39·4
39.3.2 Classic Waterfall Model 39·6
39.3.3 Rapid Application Development and Joint Application Design 39·7
39.3.4 Extreme Programming 39·8
39.3.5 Importance of Integrating Security at Every Phase 39·8

39.4 TYPES OF SOFTWARE ERRORS 39·8

39.4.1 Internal Design or Implementation Errors 39·8
39.4.2 User Interface 39·10

39.5 DESIGNING SOFTWARE TEST CASES 39·12

39.5.1 Good Tests 39·12
39.5.2 Emphasize Boundary Conditions 39·13
39.5.3 Check All State Transitions 39·14
39.5.4 Use Test-Coverage Monitors 39·14
39.5.5 Seeding 39·15
39.5.6 Building Test Data Sets 39·15
39.5.7 Fuzzing 39·16

39.6 BEFORE GOING INTO PRODUCTION 39·16

39.6.1 Regression Testing 39·16
39.6.2 Automated Testing 39·16
39.6.3 Tracking Bugs from Discovery to Removal 39·17

39.7 MANAGING CHANGE 39·17

39.7.1 Change Request 39·17
39.7.2 Tracking System 39·18
39.7.3 Regression Testing 39·18
39.7.4 Documentation 39·18

39.8 SOURCES OF BUGS AND PROBLEMS 39·19
39.8.1 Design Flaws 39·19
39.8.2 Implementation Flaws 39·19
39.8.3 Unauthorized Changes to Production Code 39·19
39.8.4 Insufficient or Substandard Programming Quality 39·19
39.8.5 Data Corruption 39·19
39.8.6 Hacking 39·20
39.9 CONCLUSION 39·21
39.10 FURTHER READING 39·21

39.1 INTRODUCTION. Software development can affect all of the six fundamental principles of information security as described in Chapter 3 in this *Handbook*, but the most frequent problems caused by poor software involve integrity, availability, and utility.

Despite the ready availability of packaged software on the open market, such software frequently has to be customized to meet its users' particular needs. Where this is not possible, programs must be developed from scratch. Unfortunately, during any software development project, despite careful planning, unforeseen problems inevitably arise. Custom software and customized packages frequently are delivered late, are faulty, and do not meet specifications. Generally, software project managers tend to underestimate the impact of technical as well as nontechnical difficulties. Because of this experience, the field of software engineering has developed; its purpose is to find reasonable answers to questions that occur during software development projects.

39.2 GOALS OF SOFTWARE QUALITY ASSURANCE. In the IEEE *Glossary of Software Engineering Terminology*, quality is defined as "the degree to which a system, component, or process meets customer or user needs or expectations." In accordance with this definition, software should be measured primarily by the degree to which user needs are met. Because software frequently needs to be adapted to changing requirements, it should be adaptable at reasonable cost. Therefore, in addition to being concerned with correctness, reliability, and usability, the customer also is concerned with testability, maintainability, portability, and compliance with the established standards and procedures for the software and development process. Software quality assurance (SQA) is an element of software engineering that tries to ensure that software meets acceptable standards of completeness and quality. SQA acts as a watchdog overseeing all quality-related activities involved in software development.

The principal goals of SQA are:

- Uncover all of a program's problems.
- Reduce the likelihood that defective programs will enter production.
- Safeguard the interests of users.
- Safeguard the interests of the software producer.

39.2.1 Uncover All of a Program's Problems. Quality must be built into a product. Testing can only reveal the presence of defects in the product. To ensure quality, SQA monitors both the development process and the behavior of software. The goal is not to pass or certify the software; the goal is to identify all the inadequacies and problems in the software so that they can be corrected.

39.2.2 Reduce the Likelihood that Defective Programs Will Enter Production. All software development must comply with the standards and policies established within the organization to identify and eliminate errors before defective programs enter production.

39.2.3 Safeguard the Interests of Users. The ultimate goal is to achieve software that meets the users' requirements and provides them with needed functionality. To meet these goals, SQA must review and audit the software development process and provide the results of those reviews and audits to management. SQA, as a functioning department, can be successful only if it has the support of management and if it reports to management at the same level as software development. Likewise, each SQA project will possess specific attributes, and the SQA program should be tailored to accommodate project needs. Characteristics to be considered include: mission criticality, schedule and budget, size and complexity of project, and size and competence of project staff organization.

39.2.4 Safeguard the Interests of Software Producers. By ensuring that software meets requirements, SQA can help prevent legal conflicts that may arise if purchased software fails to meet contractual obligations. When software is developed in-house, SQA can prevent the finger-pointing that otherwise would damage relations between software developers and corporate users.

39.3 SOFTWARE DEVELOPMENT LIFE CYCLE. Good software, whether developed in-house or bought from an external vendor, needs to be constructed using sound principles. Because software development projects often are large and have many people working on them for long periods of time, the development process needs to be monitored and controlled.

Although progress of such projects is difficult to measure, using a phased approach can control the projects. In a phased approach, a number of clearly identifiable milestones are established between the start and the finish of the project. A common analogy that is used is that of constructing a house, where the foundation is laid initially and each phase of construction is achieved in an orderly and controlled manner. Frequently, with both house construction and software development, payments for phases are dependent on reaching the designated milestones.

When developing systems, we refer to the phased process as the system development life cycle (SDLC). The SDLC is the process of developing information systems through investigation, analysis, design, coding and debugging, testing, implementation, and maintenance. These seven phases are common, although different models and techniques may contain more or fewer phases.

Generally, milestones identified in the SDLC correspond to the points in time when specified documents become available. Frequently, the documents explain and reinforce the actions taken during the just-completed phase of the SDLC. We therefore say that traditional models for the phased development are document driven to a large extent.

There are problems and drawbacks inherent in the document-driven process. The method of viewing the development process via the SDLC is not totally realistic to the actual projects. In reality, errors found in earlier phases are noted, and fixes are developed, prototyping is introduced, and solutions are implemented. This, in effect, is more than what is assumed will be necessary in the debugging and maintenance phases. Also, often the problems are solved before those later phases are reached. Because of

the recognition that much of what in traditional models is referred to as maintenance is really evolution, other models of the SDLC, known as evolutionary models, have been developed.

In traditional models, the initial development of a system is kept strictly separate from the maintenance phase. The major goal of the SDLC is to deliver a first production version of the software system to the user. It is not unusual for this approach to result in excessive maintenance costs to make the functioning or final system fit the needs of the real user. There are other models available, and it is necessary, for each project, to choose a specific SDLC model. To do this, it is necessary to identify the need for a new system and to follow this by identifying the individual steps and phases, possible interaction of the phases, the necessary deliverables, and all related materials.

There are five major system development techniques:

1. Traditional systems development life cycle (SDLC)
2. Waterfall model (a variant of the traditional SDLC)
3. Rapid application development (RAD)
4. Joint application development (JAD)
5. Extreme programming

Although these five development techniques frequently are seen as mutually exclusive, in truth they represent solutions that place different emphasis on common elements of systems design. Defined at different times, each methodology's strengths demonstrate the technology, economics, and organizational issues that were current at the time the methodology was first defined.

Software generally is constructed in phases, and tests should be conducted at the end of each phase before development continues in the next phase. Four major phases that are always present in software construction are:

1. The analysis phase, where software requirements are defined
2. The design phase, based on the previously described requirements
3. The construction, programming, or coding phase
4. The implementation phase, where software is actually installed on production hardware and finally tested before release to production

The programming phase always includes unit testing of individual modules and system testing of overall functions. Sometimes, in addition to testing during programming and implementation, a fifth phase, solely devoted to functional testing, is established. Review and modification generally follow all phases, where weaknesses and inadequacies are corrected.

39.3.1 Phases of the Traditional Software Development Life Cycle.

The seven phases discussed here comprise the most common traditional SDLC:

1. Investigation
2. Analysis
3. Design
4. Decoding and debugging

5. Testing
6. Implementation
7. Maintenance

From this traditional model, other models with fewer or more phases have been developed.

39.3.1.1 Investigation. This phase involves the determination of the need for the system. It involves determining whether a business problem or opportunity exists and conducting a feasibility study to determine the cost effectiveness of the proposed solution.

39.3.1.2 Analysis. Requirements analysis is the process of analyzing the end users' information needs, the environment of the organization and the current system, and developing functional requirements for a system to meet users' needs. This phase includes recording all of the requirements. The documentation must be referred to continually during the remainder of the system development process.

39.3.1.3 Design. The architectural design phase lists and describes all of the necessary specifications for the hardware, software, people and data resources, and information products that will satisfy the proposed system's functional requirements. The design can best be described as a blueprint for the system. It is a crucial tool in detecting and eliminating problems and errors before they are built into the system.

39.3.1.4 Coding and Debugging. This phase involves the actual creation of the system. It is done by information technology professionals, sometimes on staff within a company and sometimes from an external company that specializes in this type of work. The system is coded, and attempts are made to catch and eliminate all coding errors before the system is implemented.

39.3.1.5 Testing. Once the system is created, testing is essential. Testing proves the functionality and reliability of the system and acts as another milestone for finding problems and errors before implementation. This phase is instrumental in determining whether the system will meet users' needs. The testing should be documented in a project plan format as to the

- Objective of the testing
- Scope
- Expected results
- Test plan
- Test cases; ideally, the end users and/or business process owner(s) should be involved so that the test cases represent real-life situations as much as possible
- Test results and analysis
- Conclusion and determination of next steps
- End user/business process owner [signed] acceptance of the test results

These steps help ensure that the tests align with the users' needs and that the results are representative of what can be expected.

39.3.1.6 Implementation. Once the previous five phases have been completed and accepted, with substantiating documentation, the system is implemented and users are permitted to utilize it.

39.3.1.7 Maintenance. This phase is ongoing and takes place after the system is implemented. Because systems are subject to variances, flaws, and breakdowns, as well as difficulties when integrating with other systems and hardware, maintenance continues for as long as the system is in use.

39.3.2 Classic Waterfall Model. Although phased approaches to software development appeared in the 1960s, the waterfall model, attributed to Roy Royce, appeared in the 1970s. The waterfall model demands a sequential approach to software development and contains only five phases:

1. Requirements analysis phase
2. Design
3. Implementation
4. Testing
5. Maintenance

39.3.2.1 Analysis or Requirements Analysis. The waterfall model emphasizes analysis as part of the requirements analysis phase. Since software is always part of a larger system, requirements are first established for all system elements, and then some subset of these requirements is allocated to software. Identified requirements, for both the system and the software, are then documented in the requirements specification. A series of validation tests is required to ensure that, as the system is developed, it continues to meet these specifications.

The waterfall model includes validation and verification in each of its five phases. This means that in each phase of the software development process, it is necessary to compare the obtained results against the required results. Testing is done within every phase to answer this question and does not occur strictly in the testing phase that follows the implementation phase.

39.3.2.2 Design. The design phase is actually a multistep process focusing on data structure, software architecture, procedural detail, and interface characterization. During the design phase, requirements are translated into a representation of the software that then can be assessed for quality before actual coding begins. Once again, documentation plays an important role because the documented design becomes a part of the software configuration. Coding is incorporated into the design phase instead of being a separate phase. During coding, the design is translated into machine-readable form.

39.3.2.3 Implementation. During the implementation phase, the system is given to the user. Many developers feel that a large problem with this model is that the system is implemented before it actually is ready to be given to the user. However, waterfall model advocates claim that the consistent testing throughout the development

process permits the system to be ready by the time this phase is reached. Note that in the waterfall model, the implementation phase precedes the testing phase.

39.3.2.4 Testing. Although testing has been going on during the entire development process, the testing phase begins after code has been generated and the implementation phase has occurred. This phase focuses on both the logical internals of the software and the functional externals. Ideally, the end users and the process/business owners should be enlisted to assist in this phase, since they usually:

- Are familiar with the business requirements
- Know how the current software and hardware operates
- Know what types of activity are anticipated and planned
- Provide the ultimate acceptance approval of the software and hardware

Although not directly required by many of the regulatory reviews, involving the end users and process owners provides the IT area with additional resources, not only from a human resource perspective but also from a business knowledge and regulatory compliance perspective.

39.3.2.5 Maintenance. The maintenance phase reapplies each of the preceding life cycle steps to existing programs. Maintenance is necessary because of errors that are detected, necessary adaptation to the external environment, and functional and performance enhancements required and requested by the customer.

The waterfall model is considered to be unreliable, partly due to failure to obey the strict sequence of phases advocated by the traditional model. In addition, the waterfall model often fails because it delivers products whose usefulness is limited when requirements have changed during the development process.

39.3.3 Rapid Application Development and Joint Application Design.

Rapid application development (RAD) supports the iteration and flexibility necessary for building robust business process support. RAD emphasizes user involvement and small development teams, prototyping of software, software reuse, and automated tools. Activities must be carried out within a specified time frame known as a time box. This approach differs from other development models where the requirements are fixed first and the time frame is decided later. RAD, in an effort to keep within the time box and its immovable deadline, may sacrifice some functionality.

RAD has four phases within its cycle:

1. Requirements planning
2. User design
3. Construction
4. Cutover

The main techniques used in RAD are joint requirements planning (JRP) and joint application design (JAD). The word "joint" refers to developers and users working together through the heavy use of workshops. JAD enables the identification, definition, and implementation of information infrastructures. The JAD technique is discussed with RAD because it enhances RAD.

39.3.4 Extreme Programming. A form of *agile software development* that has gained popularity in the last two decades is *extreme programming* (XP). Programmers work closely with each other and the users to produce usable code quickly. Unit tests and integrated integration tests are constructed before the coding and constantly applied throughout the development process. A variant of the method is known as *Scrum*.

39.3.5 Importance of Integrating Security at Every Phase. Security should never be something added to software at the end of a project. Security must be considered continuously during an entire project in order to safeguard both the software and the entire system within which it functions. Regardless of which software development model is used, it is essential to integrate security within each phase of development and to include security in the testing at each phase. Rough estimates indicate that the cost of correcting an error rises tenfold with every additional phase. For example, catching an error at the analysis phase might require only several minutes—time for a user to correct the analyst—and therefore may cost only a few dollars. Catching the same error once it has been incorporated into the specifications document might cost 10 times more; after implementation, as much as 1,000 times more (this excludes any compliance-related penalties, losses, or loss responsibility).

39.4 TYPES OF SOFTWARE ERRORS

39.4.1 Internal Design or Implementation Errors. A general definition of a software error is a mismatch between a program and its specifications; a more specific definition is "the failure of a program to do what the end user reasonably expects." There are many types of software errors. Some of the most important include:

- Initialization
- Logic flow
- Calculation
- Boundary condition violations
- Parameter passing
- Race condition
- Load condition
- Resource exhaustion
- Resource, address, or program conflict with the operating system or application(s)
- Regulatory compliance considerations
- Other errors

39.4.1.1 Initialization. Initialization errors are insidious and difficult to find. The most insidious programs save initialization information to disk and fail only the first time they are used—that is, before they create the initialization file. The second time a given user activates the program, there are no further initialization errors. Thus, the bugs appear only to employees and customers when they activate a fresh copy or installation of the defective program. Other programs with initialization errors may show odd calculations or other flaws the first time they are used or initialized; because

they do not store their initialization values, these initialization errors will continue to reappear each time the program is used.

39.4.1.2 Logic Flow. Modules pass control to each other or to other programs. If execution passes to the wrong module, a logic-flow error has occurred. Examples include calling the wrong function, or branching to a subroutine that lacks a RETURN instruction, so that execution falls through the logical end of a module and begins executing some other code module.

39.4.1.3 Calculation. When a program misinterprets complicated formulas and loses precision as it calculates, it is likely that a calculation error has occurred; for example, an intermediate value may be stored in an array with 16 bits of precision when it needs 32 bits. This category of errors also includes computational errors due to incorrect algorithms.

39.4.1.4 Boundary Condition Violations. The term "boundaries" refers to the largest and smallest values with which a program can cope; for example, an array may be dimensioned with 365 values to account for days of the year and then fail in a leap year when the program increments the day-counter to 366 and thereby attempts to store a value in an illegal address. Programs that set variable ranges and memory allocation may work within the boundaries but, if incorrectly designed, may crash at or outside the boundaries. The first use of a program also can be considered a boundary condition.

39.4.1.5 Parameter Passing. Sometimes there are errors in passing data back and forth among modules. For instance, a call to a function accidentally might pass the wrong variable name so that the function acts on the wrong values. When these parameter-passing errors occur, data may be corrupted and the execution path may be affected because of incorrect results of calculations or comparisons. As a result, the latest changes to the data might be lost or execution might fall into error-handling routines even though the intended data were correct.

39.4.1.6 Race Condition. When a race occurs between event A and event B, a specific sequence of events is required for correct operation, but this sequence is not ensured by the program. For example, if process A locks resource 1 and waits for resource 2 to be unlocked while process B locks resource 2 and waits for resource 1 to be unlocked, there will be a deadly embrace that freezes the operations.

Race conditions can be expected in multiprocessing systems and interactive systems, but they can be difficult to replicate; for example, the deadly embrace just described might happen only once in 1,000 transactions if the average transaction time is short. Consequently, race conditions are among the least tested.

39.4.1.7 Load Condition. All programs and systems have limits to storage capacity, numbers of users, transactions, and throughput. Load errors can occur due to high volume, which includes a great deal of work over a long period of time, or high stress, which includes the maximum load all at one time.

39.4.1.8 Resource Exhaustion. The program's running out of high-speed memory (random access memory, or RAM), mass storage (disk), central processing unit (CPU) cycles, operating system table entries, semaphores, or other resources

can cause failure of the program. For example, inadequate main memory may cause swapping of data to disk, typically causing drastic reductions in throughput.

39.4.1.9 Interapplication Conflicts. With operating systems (OS) as complex as they are, OS manufacturers routinely distribute the code requirements and certain parameters to the application software manufacturers, so that the likelihood of program conflicts or unexpected stoppages are minimized. Although this certainly helps reduce the number of problems and improves the forward and backward compatibility with previous OS versions, on occasion even the OS vendors experience or cause difficulties when they do not conform to the parameters established for their own programs.

39.4.1.10 Other Sources of Error. It is not unusual for errors to occur where programs send bad data to devices, ignore error codes coming back, and even try to use devices that are busy or missing. The hardware might well be broken, but the software also is considered to be in error when it does not recover from such hardware conditions.

Additional errors can occur through improper builds of the executable; for example, if an old version of a module is linked to the latest version of the rest of the program, the wrong sign-on screens may pop up, the wrong copyright messages may be displayed, the wrong version numbers may appear, and various other inaccuracies may occur.

39.4.1.11 Regulatory Compliance Considerations. If an organization is subject to Sarbanes-Oxley (SOX), then documentation of the errors encountered and the resulting internal reporting and remediation efforts is critical. The error should clearly indicate how it was identified, who identified it, how it was reported to management, what the remediation will be, and when it is anticipated to be completed. Without these details, management may encounter significant difficulties in confirming that adequate internal control mechanisms exist and that it is informed and involved appropriately and adequately. Also, an error could result in a control weakness being identified by the external auditors or, even worse, as a material weakness, depending on its nature and severity.

Additionally, if and when errors are noted that affect multilocation systems or applications, the significance and materiality must be considered from the aspect of both the subsidiary and the headquarters locales. Since laws differ among states and countries, what might be considered legally acceptable standards of privacy or accounting in one locale might not be acceptable elsewhere.

39.4.2 User Interface. Generally speaking, the term "user interface" denotes all aspects of a system that are relevant to a user. It can be broadly described as the user virtual machine (UVM). This would include all screens, the mouse and keyboard, printed outputs, and all other elements with which the user interacts. A major problem arises when system designers cannot put themselves in the user's place and cannot foresee the problems that a technologically challenged user will have with an interface designed by a technologically knowledgeable person.

Documentation is a crucial part of every system. Each phase of development—requirements, analysis, development, coding, testing, errors, error solutions and modifications, implementation, and maintenance—needs to be documented. All documents and their various versions need to be retained for both future reference and auditing purposes. Additionally, it is important to document the correct use of the system and

provide adequate instructional and reference materials to the user. Security policies and related enforcement and penalties also need to be documented. Ideally, the documentation should enable any technically qualified person to repair or modify any element, as long as the system remains operational.

39.4.2.1 Functionality. A program has a functionality error if performance that can reasonably be expected is confusing, awkward, difficult, or impossible. Functionality errors often involve key features or functions that have never been implemented. Additional functionality errors exist when:

- Features are not documented.
- Required information is missing.
- A program fails to acknowledge legitimate input.
- There are factual errors or conflicting names for features.
- There is information overload.
- The material is written to an inappropriate reading level.
- The cursor disappears or is in the wrong place.
- Screen displays are wrong.
- Instructions are obscured.
- Identical functions require different operations in different screens.
- Improperly formatted input screens exist.
- Passwords or other confidential information are not obscured or protected adequately.
- Tracing the user data entry or changes is unavailable or incomplete.
- Segregation of duties is not enforced. (This can be particularly critical for organizations subject to the Health Insurance Portability and Accountability Act [HIPAA], SOX, International Organization for Standardization [ISO] 17799, and/or the Gramm-Leach-Bliley Act [GLBA].)

39.4.2.2 Control (Command) Structure. Control structure errors can cause serious problems because they can result in:

- Users getting lost in a program
- Users wasting time because they must deal with confusing commands
- Loss of data or the unwanted exposure of data
- Work delay
- Financial cost
- Unanticipated exposure to data leakage or compromise; this can result in significant liability if consumers' personal identifying information (PII) is compromised
- Data not being encrypted as intended or being visible to unauthorized users

Some common errors include:

- Inability to move between menus
- Confusing and repetitive menus

- Failure to allow adequate command-line entries
- Requiring command-line entries that are neither intuitive nor clearly defined on screen
- Failure of the application program to follow the operating system's conventions
- Failure to distinguish between source and parameter files, resulting in the wrong values being made available to the user through the interface, or failure to identify the source of the error
- Inappropriate use of the keyboard when new programs do not meet the standard of a keyboard that has labeled function keys tied to standard meanings
- Missing commands from the code and screens, resulting in the user being unable to access information, to utilize programs, or to provide for the system to be backed up and recoverable, as well as a host of other commands that leave the system in a state of less-than-optimum operability
- Inadequate privacy or security that can result in confidential information being divulged, the complete change or loss of data without recoverability, poor reporting, and even undesired access by outside parties

39.4.2.3 Performance. Speed is important in interactive software. If a user feels that the program is working slowly, that can be an immediate problem. Slow operation can depend on (but is not limited to) the OS, the other applications running, memory allocation, memory leakage, and program conflicts. At another level, performance suffers when program designs make it difficult to change their functionality in response to changing requirements. Performance errors include slow response, unannounced case sensitivity, uncontrollable and excessively frequent automatic saves, inability to save, and limited scrolling speed.

39.4.2.4 Output Format. Output format errors can be frustrating and time consuming. An error is considered to have occurred when the user cannot change fonts, underlining, boldface, and spacing that influence the final look of the output; alternatively, delays or errors when printing or saving document may occur. Errors occur when the user cannot control the content, scaling, and look of tables, figures, and graphs. Additionally, there are output errors that involve expression of the data to an inappropriate level of precision.

39.5 DESIGNING SOFTWARE TEST CASES

39.5.1 Good Tests. No software program can ever be tested completely, since it would not be cost effective or efficient to test the validity, parameters, syntax, and boundaries of every line of code. Even if all valid inputs are defined and tested, there is no way to test all invalid inputs and all the variations on input timing. It is also difficult, if not impossible, to test every path the program might take, find every design error, and prove programs to be logically correct. Nevertheless, a good test procedure will find most of the problems that would occur, allowing the designers and developers to correct those problems and ensure that the software works properly. Generally, OS and application manufacturers classify the severity level of errors encountered, and many (as evidenced by the frequency and content of program patches) correct only the most serious ones that either they or the users notice.

Over the past several years, there has been an increasing debate on whether software vulnerabilities should be publicized immediately by researchers, or if the researchers should automatically notify and give the manufacturer time to correct the flaw. Generally, from a casual review of the vulnerabilities discovered concerning Microsoft's Internet Explorer version 6 over the past few years, the trend appears to have been that the researcher notifies Microsoft and allows Microsoft a varying period of time to create and distribute the patch (e.g., three to six months) before the researcher publicly announces the weakness. This trend is based on published reports in the technical news media.

Whenever you expect the same results from two tests, you consider the tests equivalent. A group of tests forms an equivalence class if the tester believes that the tests all test the same thing, and if one test catches or does not catch a bug, the others probably will do the same. Classical boundary tests check a program's response to input and output data; equivalence tests, however, teach a way of thinking about analyzing programs that enhances and strengthens test planning.

Finding equivalence classes is a subjective process. Different people analyzing the same program will come up with different lists of equivalence classes because of what the programs appear to achieve. Test cases often are lumped into the same equivalence class when they involve the same input variables, result in similar operations, affect the same output variables, or handle errors in the same manner.

Equivalence classes can be groups of tests dealing with ranges or multiple ranges of numbers, members of a group, and even time determined. Equivalence classes are generally the most extreme values, such as the biggest, smallest, fastest, and slowest. When testing, it is important to test each edge of an equivalence class, on all sides of each edge. Testers should use only one or two test cases from each equivalence class because a program that passes the tests generally will pass any test drawn from that class. Invalid input equivalence classes often allow program bugs to be overlooked during debugging.

39.5.2 Emphasize Boundary Conditions. Boundaries are crucial for checking each program's response to input and output data. Equivalence class boundaries are the most extreme values of the class; for example, boundary conditions may consist of the biggest and smallest, soonest and latest, shortest and longest, or slowest and fastest members of an equivalence class.

Tests should include values below, at, and above boundary values. Additionally, they should be tested at various user levels (e.g., administrator, root or super user, data entry, and read-only access). This testing must ensure that there are no conditions at any level that permit unintended access or unauthorized escalation of privileges.

When programs fail with nonboundary values, they generally fail at the boundaries too. Programs passing these tests probably also will pass any other test drawn from that class. Tests should also be able to generate the largest and smallest legitimate output values, remembering that input-boundary values may not generate output-boundary values.

Today, the most prevalent security breaches involve buffer overflows in active code (ActiveX and Java), in which inadequate bounds checking on input strings allows overflowing text to be interpreted as code that then can carry out improper operations. Restrictions on input length and type would prevent these exploits. Such overflows can result in the exposure of consumers' PII and incurring unwanted liability for the exposure; currently, the liability exists principally at the state level and may include criminal sanctions, depending on the state in which the consumers reside.

39.5.3 Check All State Transitions. All interactive programs move from one state to another. A program's state is changed whenever something causes the program to alter output (e.g., to display something different on the screen) or to change the range of available choices (e.g., displaying a new menu of user options).

To test state transitions, the test designer should lay out a transition probability matrix to show all the paths people are likely to follow. State transitions can be complex and often depend not on the simple choices provided but on the numbers the user enters. Testers often find it useful to construct menu maps that show exactly where to go from each choice available.

Menu maps also can show when and where users go when menu or keyboard commands take them to different states or dialogs. Maps are particularly handy when working with spaghetti code (code that has been poorly designed or badly maintained so that logical relationships are difficult to see); maps allow the designer or user to reach a dialog box in several ways and then proceed from the dialog box to several places. For spaghetti code, the menu maps afford a simpler method of spotting relationships between states than trying to work exclusively from the program itself because the map shows transition between states on paper or on screen. After mapping the relationships, the designer or user can check the program against the map for correctness.

Full testing theoretically requires that all possible paths are tested. However, in practice, complete testing may be unattainable for complex programs. Therefore, it makes sense to test the most frequently used paths first.

39.5.3.1 Test Every Limit. It is necessary to test every limit on a program's behavior that is specified by any of the program's documents. Limits include the size of files, the number of terminals, the memory size the program can manage, the maximum size it requires, and the number of printers the program can drive. It is important to check on the ability of the program to handle large numbers of open files on an immediate basis and on a long-term, continuing basis as well. Load testing is actually boundary condition testing and should include running tests the program ought to pass and tests the program should not pass.

For Web applications, tools exist to simplify this task, as some will check the boundaries and limits of fields at varying depths of data entry to test the protections against buffer overflow, invalid data, and incorrect data syntax; such tools include but are not limited to WebInspect by Spi Dynamics.

39.5.3.2 Test for Race Conditions. After testing the system under "normal" load, it should be tested for race conditions. A "race condition" is usually defined as anomalous behavior due to unexpected critical dependence on the relative timing of events. Race conditions generally involve one or more processes accessing a shared resource such as a file or variable, where this multiple access has not been properly controlled. Systems that are vulnerable to races, especially multiuser systems with concurrent access to resources, should undergo a full cycle of testing under load. Race conditions sometimes can be identified through testing under heavy load, light load, fast speed, slow speed, multiprocessors running concurrent programs, enhanced and more numerous input/output devices, frequent interrupts, less memory, slower memory, and related variables.

39.5.4 Use Test-Coverage Monitors. For complex programs, path testing usually cannot or would not test every possible path throughout a program for practical

as well as cost-effectiveness reasons. A more practical glass box approach (i.e., with knowledge of the internal design of a program) is to use the source code listing to force the program down every branch visible in the code. When programmers add special debugging code during development, a unique message will print out or be added to a log file whenever the program reaches a specified point. Typically, such messages can be generated by calling a print routine with a unique parameter for each block of code. When source code is compiled, many languages permit a switch to be set to allow conditional compilation, so that a test version of the code contains active debugging statements whereas a production version does not. For interpreted code, such as most fourth-generation languages, similar switches allow for inclusion or activation of debugging instructions.

Since distinct messages are planted, it is possible to know exactly what point in the program the test has reached. Programmers specifically insert these messages at significant parts of the program. Once these messages have been added, anyone can run the program and conclude whether the different parts actually have run and been tested.

Special devices or tools also can be used to add these messages to the code automatically. Once source code is fed to such a coverage monitor, it analyzes the control structures in the code and adds probes for each branch of the program. Adding these probes or lines of code is called instrumenting the program. It is possible to tell that the program has been pushed down every branch when all the probes have been printed. Besides instrumenting code, a good coverage monitor can capture probe outputs, perform analyses, and summarize them. Some coverage monitors also log the time used for each routine, thus supporting performance analysis and optimization.

A coverage monitor is designed for glass box testing, but knowledge of the internals of the program under test is not needed in order to use the monitor. The coverage monitor counts the number of probe messages, reports on the number of probes triggered, and even reports on the thoroughness of the testing. Because it is possible for the coverage monitor to report on untriggered branches, the monitor can find and identify code that does not belong, such as routines included by default but never used or deliberately inserted undocumented code (a Trojan horse). Some commercially available coverage monitors are, unfortunately, themselves full of bugs. It is important to obtain a full list of the known bugs and patches from the developer. In addition, such tools should themselves be tested with programs that contain known errors, in the process known as seeding, to verify their correctness.

39.5.5 Seeding. Quality assurance procedures themselves benefit from testing. One productive method, known as seeding, is to add known bugs to a program and measure how many of them are discovered through normal testing procedures. The success rate in identifying such known bugs can help estimate the proportion of unknown bugs left in a program. Such seeding is particularly important when establishing automated testing procedures and test-coverage monitors; also, it is useful for SOX or other compliance-related testing.

39.5.6 Building Test Data Sets. One of the most serious errors in quality assurance is to allow programmers to use production data sets for testing. Production data may include confidential information that should not be accessible to programming staff, so access should be forbidden not only to production data but even to copies of production data, or even to copies of subsets or samples from production data. However, *anonymizing* processes applied to copies of production data (e.g., scrambling names

and addresses of patients in a medical records database) may produce test data sets suitable for use in quality assurance.

Alternatively, if the organization has a test system that is completely (logically and physically) segregated from the production system, using historical data may be appropriate. In doing so, particular care and appropriate documentation must be maintained to ensure that there is no chance of the test system interfacing with the production system. A side benefit of such testing is that it may help the organization's SOX or HIPAA periodic testing and documentation regarding the availability, utility, and integrity of backup media and systems.

39.5.7 Fuzzing. A useful addition to the techniques described above is the automated injection of random data into the input stream. *Fuzzing* can support the planned challenges to a system with unexpected inputs that programmers may not have thought of.

39.6 BEFORE GOING INTO PRODUCTION

39.6.1 Regression Testing. Regression testing is fundamental work done by glass box and black box testers. The term is used in two different ways. The first definition involves those tests where an error is found and fixed and the test that exposed the problem is performed again. The second definition involves finding and fixing an error and then performing a standard series of tests to make certain the changes or fix made did not disturb anything else.

The first type of regression testing serves to test that a fix does what it is intended to do. The second type tests the fix and also tests the overall integrity of the program. It is not unusual for people who mention regression testing to be referring to both definitions, since both involve fixing and then retesting. It is recommended that both types of testing be done whenever errors are fixed. If an organization is subject to SOX or HIPAA, appropriate documentation should be maintained that details the test methodology, the sample selection, the frequency, the results, and the remediation (if any).

39.6.2 Automated Testing. In some enterprises, every time a bug is fixed, every time a program is modified, and every time a new version is produced, a tester runs a regression test. All of this testing takes a considerable amount of time and consumes both personnel and machine resources. In addition, repetitive work can become mind-numbing, so testers may accidentally omit test cases or overlook erroneous results.

To cut down on the time consumption of personnel and the repetitive nature of the task, it is possible to program the computer to run acceptance and regression tests. This type of automation results in execution of the tests, collection of the results, comparison of the results with known good results, and a report of the results to the tester.

Early test automation consisted essentially of keyboard macros or scripts, with limited capacity for responding to errors; typically, an error would produce invalid results that would be fed into the next step of processing and lead to long chains of meaningless results. Slightly more advanced test harnesses would halt the test process at each error and require human intervention to resume. Both methods were of only modest help to the test team. However, today's test-automation software can include test databases that allow orderly configuration of tests, specific instructions for restarting test sequences after errors, and full documentation of inputs and results. For realistic load testing, some systems can be configured to simulate users on workstations, with

scripts that define a range of randomly generated values, specific parameter limits, and even variable response times. Large-scale load testing can connect computers together through networks to simulate thousands of concurrent users and thus speed identification of resource exhaustion or race conditions.

Test automation can be well worth the expenditures required. Automated testing is usually more precise, more complete, faster, and less expensive than the tests done by human personnel. Typically, automated tests can accomplish tenfold or hundredfold increases in the number of tests achievable through manual methods. In addition to freeing personnel to do more rewarding work, such methods greatly reduce overall maintenance costs due to detection of errors before systems reach production. Additionally, for regulatory purposes, the automated testing requires significantly less sampling, documentation, and testing. Appropriate management review and reporting of results still are required, as is documentation of remediation follow-up.

39.6.3 Tracking Bugs from Discovery to Removal. Finding a bug is not enough. Even eliminating a bug is insufficient. A system also must allow the causes of each bug to be identified, documented, and rectified. For these reasons, it is important to track and document all problems from the time of their discovery to their removal, to be certain that the problems have been resolved and will not affect the system. Too, this provides important historical records in diagnosing and remediating incidents and provides appropriate documentation for audit and regulatory purposes.

Problem-tracking systems must be used to report bugs, track solutions, and write summary reports about them. An organized system is essential to ensure accountability and communication regarding the bugs. Typical reports include: where bugs originate (e.g., which programmers and which teams are responsible for the greatest number of bugs); types of problems encountered (e.g., typographical errors, logic errors, boundary violations); and time to repair. However, problem-tracking systems can raise political issues, such as project accountability, personal monitoring, control issues, and issue remediation regarding the data in the database and who owns them.

Once a tracking system is established, a bug report is entered into the database system and a copy goes to the project manager, who either prioritizes it and passes it along or responds to it personally. Eventually, the programmers will be brought into the loop, will fix the problem, and will mark the problem as fixed. The fixed problem then is tested and a status report is issued. If a fix proves not to work, that becomes a new problem that needs to be addressed.

39.7 MANAGING CHANGE. Many people may be involved in creating and managing systems; whenever changes are made to the system, those changes need to be managed and monitored in an organized fashion to avoid chaos.

39.7.1 Change Request. The change request is an important document that requests either a fix or some other change to a program or system. Information contained on the change request form generally includes who is requesting the change, the date the request is being made, the program and area affected, the date by which the change is needed, and authorization to go ahead and make the change.

Most companies have paper forms that are retained for monitoring and auditing purposes. As the popularity of digital signatures grows and society continues to move toward a paperless office, it seems logical that some of these change requests eventually will become totally automated.

39.7.2 Tracking System. The tracking system may be manual, automated, or a combination of methods; for regulatory and control purposes, a system that incorporates automated tracking is best, since it reduces the chance for error, and it can enhance follow-up and remediation. Regardless of how it is kept, the original change form generally is logged into a system, either manually filed or entered into an automated database by some identifying aspect, such as date, type of change, or requesting department.

The system is used to track what happens to the change. When an action is taken, information must be entered into the system to show who worked on the request, what was done, when action was taken, what the result of the action was, who was notified of the actions taken, whether the change was accepted, and related information. This information is crucial in ensuring that issues are managed appropriately, escalated as needed to senior management, and tracked until remediated.

39.7.3 Regression Testing. When a change or fix has been made, regression testing verifies that the change has indeed been made and that the program now works in the fashion desired, including all other functions. The successful completion of the testing should be documented in writing; this can be through the tracking system or in documentation appended to the tracking system entry.

39.7.4 Documentation. Documentation is crucial when considering, approving, and implementing changes. If information is retained strictly in an individual's head, what happens when the individual goes on vacation, is out sick, or leaves the enterprise permanently? What if the individual simply does not remember what changes were made, or when; how does an organization know who was involved, what actually was done, and who signed off that the changes were made and accepted?

Undocumented changes can result in:

- System crashes
- Inconsistent data entry
- Inappropriate segregation of duties
- Data theft or corruption
- Embezzlement
- Other serious crimes

Lack of documentation can mean that unauthorized changes were made and likely can result in regulatory and audit violations. Undocumented changes may violate segregation of duties (e.g., a person might be informally approving changes to his or her own account, approving changes carried out by a supervisor, or authorizing changes by a person who is not under the responsibility of the approver).

Lack of documentation is often a violation of corporate policy and is often cited in audits; it often will cause significant difficulties regarding SOX, GLBA, and HIPAA compliance efforts and reviews, since management's assertions might not be supported adequately. Because constructing adequate documentation after the fact is almost impossible (and frequently can be considered to be falsification of records in sectors such as financial services), documentation must proceed in step with every phase of a system's development and operation. Thereafter, documentation must be retained according to company policy and legal requirements.

39.8 SOURCES OF BUGS AND PROBLEMS. Locating and fixing bugs is an important task, but it is also essential to try to determine where and why the bugs and related problems originated. By finding and studying the source, often it is possible to prevent a recurrence of the same or similar type of problem.

39.8.1 Design Flaws. Design flaws often occur because of poor communication between users and designers. Users are not always clear about what they want and need, while designers misunderstand, misinterpret, or simply ignore what the users relate to them. Within the design process, even when users' feelings and requirements are known and understood, design flaws can occur. However, they are easier to identify and remedy if the appropriate documentation is affected. Without the proper documentation, it may not be possible to identify the source of an error; this can result in patches or remedies that need to be patched. Often flaws result from attempts to comply with unrealistic delivery schedules. Managers should support their staff in resisting the pressure to rush through any of the design, development, and implementation stages.

39.8.2 Implementation Flaws. Whenever a program is developed and implemented, there are time limits and deliverable dates. Most problems during development cause delay, so developers and testers often are rushed to get the job done. Sometimes they sacrifice the documentation, the review portion of the project, or even cut short testing, leaving unrecognized design and implementation flaws in place. Managers should emphasize the value of thorough review and testing, and allocate enough time to avoid such blunders.

39.8.3 Unauthorized Changes to Production Code. If problems are traced to unauthorized changes, project managers should examine their policies. If the policies are clear, perhaps employee training and awareness programs need improvement. Managers also should examine their own behavior to ensure that they are not putting such pressure on their staff that cutting corners is perceived to be acceptable.

39.8.4 Insufficient or Substandard Programming Quality. A programmer can make or break a program and a project. Programmers play essential roles in software development. Therefore, it is important that all programmers on a project be capable and reliable. Project programmers should be carefully screened before being placed on a software development project to ensure that their skills meet project requirements. Sometimes holes in programmers' skills can be filled with appropriate training and coaching. However, if problems appear to be consistent, management may want to check a programmer's background to verify that he or she did not falsify information when applying for the job. Truly incompetent programmers need to be removed from the project. Additionally, if management suspects that the programming quality is inadequate, then it should consider having an independent programmer review the code, performing particularly rigorous quality assurance (QA) testing.

39.8.5 Data Corruption. Data corruption can occur because of poor programming, invalid data entry, inadequate locking during concurrent data access and modification, illegal access by one process to another process data stack, and hardware failures. Data corruption can occur even when a program is automatically tested or run without human intervention. In any event, when searching for the sources of bugs and other problems, it is important to do a careful review of the data after each round of testing, in order to identify deviations from the correct end state. The review should

be documented appropriately and escalated to management; it then will assist the regulatory compliance efforts.

39.8.6 Hacking. When bugs and problems do occur, hacking—both internal and external—should be considered a possibility and searched for by reviewing the logs of who worked on the software and when that work was done. Managers should be able to spot the use of legitimate IDs when the individuals involved were on vacation, out sick, or not available to log on for other reasons.

Archiving logs and retrieving them is an increasingly critical capability, aside from the regulatory requirements of SOX, HIPAA, GLBA, and so on. Logs are often used for:

- Incident response activities, particularly in determining:
  - If an incident occurred
  - When an incident occurred
  - What happened prior to and subsequent to when an incident is suspected to have occurred
- Research; this can be to review performance capabilities or issues of hardware or software or networking

To ensure that the logs are available when needed, the organization needs a sound strategy that preserves and archives logs periodically, so that an attacker or a system anomaly does not wipe them from the system's hard drive. Although there are many means to accomplish this, a frequently used inexpensive method involves establishing a log server. For example, an organization might identify a discarded desktop computer, outfit it with a new, high-speed, high-capacity optical read-only memory (ROM) drive, and increase memory capacity (if needed). By routing the logs to the optical drive frequently and recording them throughout the day (perhaps multiple times during an hour) and changing the optical disk at least daily, the organization will create a forensically sound archive that can be duplicated easily for research, law enforcement, or legal use. Often the daily logs volume will not fill a DVD or a CD. However, having the regular routine of changing the media daily helps ensure that the media does not reach its capacity at an inopportune time—indeed, if a worm such as Blaster or Nimda infects the organization, then the logs likely will swell rapidly. Too, if an attacker within the system stops the logging, the organization should encounter less difficulty determining the date, time (or time period), and extent of the intrusion because ROM cannot be overwritten, in contrast with read-write (RW) media.

Unauthorized changes to code and data sometimes can be identified even if the perpetrator stops the logging and deletes or modifies log files. One method is to create checksums for production or other official versions of software and to protect the checksums against unauthorized access and modification using encryption and digital signatures. Similarly, unauthorized data modifications sometimes can be made more difficult by creating checksums for records and linking the checksums to time and data stamps and to the checksums for authorized programs. Under those conditions, unauthorized personnel and intruders find it difficult to create valid checksums for modified records. Other products, such as intrusion detection systems (IDSs) or intrusion prevention systems (IPSs), may be useful in an organization's protection and risk-management strategy.

39.9 CONCLUSION. This chapter has presented a comprehensive overview of the software development and quality assurance processes that must be utilized when developing, implementing, and modifying software. Software development involves more than simply selecting and utilizing an approach such as the traditional, waterfall, or RAD methodology. It means working as a team to develop, review, refine, and implement a viable working product. Many good techniques and products can be applied to both the SDLC and the quality assurance portions of producing software; some of the key elements are good documentation, allowing sufficient time for testing in the development and maintenance processes, building good tests, establishing test data, automating testing, and keeping track of change requests.

39.10 FURTHER READING

Basandra, S. *Software Architecture, Data Structures, Algorithms, Programming and Testing Questions and Answers*. Basandra Books, 2013.

Black, D. B. *Software Quality Assurance (Building Better Software Better)*. Amazon Digital Publishing, 2012.

Chemuturi, M. *Mastering Software Quality Assurance: Best Practices, Tools and Techniques for Software Developers*. J. Ross Publishing, 2010.

Fox, C., and P. Zonneveld. *IT Control Objectives for Sarbanes-Oxley: The Role of IT in the Design and Implementation of Internal Control Over Financial Reporting*, 2nd ed. IT Governance Institute, 2006.

Godbole, N. S. *Software Quality Assurance: Principles and Practices*, 2nd ed. Alpha Science International Ltd., 2013.

Hambling, B., and P van Goethem. *User Acceptance Testing—A Step-by-Step Guide*. BCS, The Chartered Institute for IT, 2013.

Horch, J. W. *Software Quality Assurance: A Standards-based Guide*. IEEE Computer Society, 2013.

ISO/IEC. "ISO/IEC 29119 Software Testing: The new international software testing standard." ISO/IEC, 2013. www.softwaretestingstandard.org/index.php

McCaffrey, J. D. *Software Testing: Fundamental Principles and Essential Knowledge*. BookSurge Publishing, 2009.

OWASP. "Fuzzing." Open Web Application Security Project, 2013. https://www.owasp.org/index.php/Fuzzing

Rasmussen, J. *The Agile Samurai: How Agile Masters Deliver Great Software*. Pragmatic Bookshelf, 2013.

Rubin, K. S. *Essential Scrum: A Practical Guide to the Most Popular Agile Process*. Addison-Wesley Professional, 2012.

Shaffer, S. C. *A Brief Introduction to Software Development and Quality Assurance Management*. S. C. Shaffer Publications, 2013.

Tassey, G. *The Economic Impacts of Inadequate Infrastructure for Software Testing: Final Report.* RTI Project Number 7007.011. NIST, 2002.

Wells, D. *Extreme Programming: A Gentle Introduction.* XP, 2009, www.extremeprogramming.org

Wood, D. C. Principles of Quality Costs: Financial Measures for Strategic Implementation of Quality Management, 4th ed. Quality Press, 2013.

CHAPTER 40

MANAGING SOFTWARE PATCHES AND VULNERABILITIES

Karen Scarfone, Peter Mell, and Murugiah Souppaya

40.1 INTRODUCTION 40·1

40.2 THE IMPORTANCE OF PATCH MANAGEMENT 40·2

40.3 THE CHALLENGES OF PATCH MANAGEMENT 40·2

- 40.3.1 Timing, Prioritization, and Testing 40·3
- 40.3.2 Patch Management Configuration 40·4
- 40.3.3 Alternative Host Architectures 40·5
- 40.3.4 Other Challenges 40·6

40.4 ENTERPRISE PATCH MANAGEMENT TECHNOLOGIES 40·7

- 40.4.1 Components and Architecture 40·7
- 40.4.2 Security Capabilities 40·8
- 40.4.3 Management Capabilities 40·9

40.5 METRICS AND MEASURES 40·10

40.6 FURTHER READING 40·11

40.7 NOTES 40·11

40.1 INTRODUCTION. *Vulnerabilities* are flaws that can be exploited by a malicious entity to gain greater access or privileges than it is authorized to have on a computer system. *Patches* are additional pieces of code developed to address problems (commonly called "bugs") in software. Patches correct security and functionality problems in software and firmware. *Patch management* is the process for identifying, acquiring, installing, and verifying patches for products and systems. Patch management is a security practice designed to proactively prevent the exploitation of IT vulnerabilities that exist within an organization. The expected result is to reduce the time and money spent dealing with vulnerabilities and the exploitation of those vulnerabilities. Proactively managing vulnerabilities of systems will reduce or eliminate the potential for exploitation and involve considerably less time and effort than responding after an exploitation has occurred.

This chapter is designed to assist organizations in implementing patch and vulnerability remediation programs. It first explains the importance of patch management and then discusses the challenges inherent in performing patch management. Next, the chapter provides an overview of enterprise patch management technologies and gives recommendations for their use. It also seeks to inform the reader about possible measures and metrics for enterprise patch management.[1]

For more general information about related topics, see selected chapters in this *Handbook*:

Chapter 38: Writing Secure Code
Chapter 39: Software Development and Quality Assurance
Chapter 46: Vulnerability Assessment
Chapter 47: Operations Security and Production Controls
Chapter 52: Application Controls
Chapter 53: Monitoring and Control Systems

40.2 THE IMPORTANCE OF PATCH MANAGEMENT. From a security perspective, patches are most often of interest because they are mitigating software flaw vulnerabilities; applying patches to eliminate these vulnerabilities significantly reduces the opportunities for exploitation. Also, patches are usually the most effective way to mitigate software flaw vulnerabilities, and are often the only fully effective solution. Sometimes there are alternatives to patches, such as temporary workarounds involving software or security control reconfiguration, but these workarounds often negatively impact functionality.

Patches serve other purposes than just fixing software flaws; they can also add new features to software and firmware, including security capabilities. New features can also be added through upgrades, which bring software or firmware to a newer version in a much broader change than just applying a patch. Upgrades may also fix security and functionality problems in previous versions of software and firmware. Also, vendors often stop supporting older versions of their products, which includes no longer releasing patches to address new vulnerabilities, thus making older versions less secure over time. Upgrades are necessary to get such products to a supported version that is patched.

As Section 40.3 of this chapter explains, there are several challenges that complicate patch management. If organizations do not overcome these challenges, they will be unable to patch systems effectively and efficiently, leading to easily preventable compromises. Organizations that can minimize the time they spend dealing with patching can use those resources for addressing other security concerns. Already many organizations have largely operationalized their patch management, making it more of a core IT function than a part of security. However, it is still important for all organizations to carefully consider patch management in the context of security because patch management is so important to achieving and maintaining sound security.

Patch management is required by various security compliance frameworks, mandates, and other policies. For example, NIST Special Publication (SP) 800-53[2] requires the SI-2, Flaw Remediation security control, which includes installing security-relevant software and firmware patches, testing patches before installing them, and incorporating patches into the organization's configuration management processes. Another example is the Payment Card Industry (PCI) Data Security Standard (DSS),[3] which requires that the latest patches be installed and sets a maximum time frame for installing the most critical patches.

40.3 THE CHALLENGES OF PATCH MANAGEMENT. This section briefly examines the challenges inherent in performing patch management. These are the challenges that the patch management technologies discussed in Section 40.4 of this chapter are trying to solve.

40.3.1 Timing, Prioritization, and Testing. Timing, prioritization, and testing are intertwined issues for enterprise patch management. Ideally, an organization would deploy every new patch immediately to minimize the time that systems are vulnerable. However, in reality this is simply not possible because organizations have limited resources, which makes it necessary to prioritize which patches should be installed before other patches. Further complicating this is the significant risk of installing patches without first testing them, which could cause serious operational disruptions, potentially even more damaging than the corresponding security impact of not pushing the patches out. Unfortunately, testing patches consumes even more of the limited resources and makes prioritization even more important. For patch management, timing, prioritization, and testing are often in conflict.

Product vendors have responded to this conflict by bundling patches for their products. Instead of releasing dozens of patches one at a time over a period of three months, necessitating testing and patch deployment every few days, a vendor might release their patches in a single bundle once a quarter. This allows an organization to perform testing once and roll out patches once, which is far more efficient than testing and rolling out all the patches separately. It also reduces the need to prioritize patches—the organization just needs to prioritize the bundle instead of separately prioritizing each patch it contains. Vendors who bundle patches tend to release them monthly or quarterly, except for cases when an unpatched vulnerability is actively being exploited, in which case they usually issue the appropriate patch immediately instead of delaying it for the next bundle.

There is a downside to patch bundling; it lengthens the time from the discovery of a vulnerability to the time a patch for it becomes publicly available. If an attacker discovers the same vulnerability before the patch is released, the attacker may have a longer window of opportunity to exploit the vulnerability because of the intentional delay in releasing the patch. However, there are two mitigating factors here. One is that if exploitation is known to be occurring, the vendor is likely to release the patch immediately. The other factor is that patches may be installed more quickly if they are bundled than if they are all released separately. So that effectively helps to shrink the window of opportunity for vulnerabilities associated with bundled patches.

There are even more issues to consider with timing. The release of a patch may provide attackers with the information that they need to exploit the corresponding vulnerability (e.g., reverse engineer the vulnerability from the patch), meaning that a newly released patch might need to be applied immediately to avoid compromises. However, if a vulnerability is not being exploited yet, organizations should carefully weigh the security risks of not patching with the operational risks of patching without performing thorough testing first. In some operational environments, such as virtual hosts with snapshot capabilities enabled, it may be preferable to patch without testing as long as the organization is fully prepared to roll back the patches if there are usability or functionality problems caused by them.

Another fundamental issue with timing is forcing the implementation of changes to make a patch take effect; this can require restarting a patched application or service, rebooting the operating system, or making other changes to the state of the host. Ultimately what matters is not when the patch was installed, but when the patch actually takes effect. In some cases it may make more sense to mitigate a vulnerability through an alternative method, at least until patches are fully operational. An example is changing configuration settings for vulnerable software to temporarily block vulnerable application functionality. Each mitigation option has different implications for the security, functionality, and operations of the vulnerable host, so it is not a trivial

matter to select one option over others. Also, if configuration settings are changed, this necessitates preserving the old setting values and restoring them at the appropriate time. Another problem with changing configuration settings is that they often require a state change to the host to take effect, such as restarting an application. Implementing configuration changes may be as disruptive to the operations of a host as installing a patch.

> Forcing the implementation of changes that requires rebooting the operating system can be problematic when the host requires authentication before booting, such as the use of full disk encryption (FDE) software. Organizations using FDE software or other technologies that require authentication before booting should carefully consider the impact that these technologies may have on patch installation.

Prioritizing which patches to apply and when to apply them is closely related to timing, but there are other considerations as well. It can depend on the relative importance of the vulnerable systems (for example, servers versus clients) and the relative severity of each vulnerability (e.g., vulnerability severity metrics such as the Common Vulnerability Scoring System [CVSS]). Another consideration is dependencies that patches may have on each other; installing one patch may require installing other patches first, and in some cases restarting an application or rebooting a host multiple times to make the patches take effect sequentially.

In summary, organizations should carefully consider the relevant issues related to timing, prioritization, and testing when planning and executing their enterprise patch management processes.

40.3.2 Patch Management Configuration. Another major challenge in enterprise patch management is that there are usually multiple mechanisms for applying patches. For example:

- A piece of software may be able to automatically update itself.
- A centralized OS management tool may be able to initiate patching.
- Third-party patch management applications may be able to initiate patching.
- Network access control, health check technologies, and similar technologies may be able to initiate patching.
- A user may be able to manually direct software to update itself.
- A user may be able to manually install a patch or a new version of the software.

Having multiple ways of applying patches can cause conflicts. Multiple methods might each try to patch the same software, which is particularly problematic when the organization doesn't want certain patches applied because of issues with those patches, testing delays, etc. Multiple methods can also cause patches to be delayed or missed, because each tool or administrator may assume another one is already taking care of a particular patch. Organizations should identify all the ways in which patches could be applied and act to resolve any conflicts among patch application methods.

A related problem with patch management configuration is that users may override or circumvent patch management processes. If users are able to make changes to their hosts' software, such as altering settings (e.g., enabling direct updates, disabling patch

management software), installing old versions of software, and uninstalling patches, they can undermine patch management integrity. To address these problems, organizations should ensure that users cannot disable or otherwise negatively affect enterprise patch management technologies, and organizations should perform continuous monitoring of enterprise patch management technologies to identify any issues that occur.

40.3.3 Alternative Host Architectures. Enterprise patch management is relatively straightforward when all of the hosts are fully managed and running typical applications and operating systems on a regular platform. When alternative host architectures are employed, patch management can be considerably more challenging. Examples of these architectures include the following:

Unmanaged hosts. As discussed in Section 40.3.2 of this chapter, it can be much more difficult to control patching when hosts are not centrally managed (i.e., users manage their own hosts).

Out-of-office hosts (e.g., telework laptops). Hosts on other networks are not protected by the enterprise's network security controls (firewalls, network intrusion detection systems, vulnerability scanners, etc.).

Nonstandard IT components (e.g., appliances). On such hosts, it's often not possible to patch individual applications independently. Rather, the organization must wait for the component vendor to release updated software.

Mobile devices. Smartphones, tablets, and other mobile devices (excluding laptops) typically run mobile operating systems, and patching for these devices is fundamentally different. It is often necessary to connect the mobile device to a desktop or laptop and to acquire and download updates through that desktop or laptop. Some mobile devices can directly download updates, but this can be problematic because of bandwidth considerations (such as taking a long time to download large updates and paying data charges for the downloads). Another option for keeping mobile devices updated is the use of enterprise mobile device management software. Enterprise mobile device management software is used to manage mobile devices, even personally owned devices not controlled by the organization. It can install, update, and remove applications, and it can restrict enterprise access if the phone's operating system and mobile device management software are not up to date. See Section 3 of SP 800-124 Revision 1, Guidelines for Managing and Securing Mobile Devices in the Enterprise, for more information.

Operating system (OS) virtualization. Patches need to be maintained for every OS image and snapshot used for full virtualization. Patching capabilities are often built into virtualized environments, such as the ability to patch offline images and quarantine dormant virtual machine instances. See NIST SP 800-125, Guide to Security for Full Virtualization Technologies, for additional information—specifically, Section 3.3 discusses virtual machine image and snapshot management.

Firmware. Firmware updates, such as updating the system BIOS, generally require special privileges and involve different procedures than other types of updates. See NIST SP 800-147, BIOS Protection Guidelines, for additional information on BIOS updates.

Organizations should carefully consider all alternative host architectures in use for the enterprise when designing enterprise patch management policies and solutions.

40.3.4 Other Challenges. This section briefly discusses other challenges not covered earlier in this section.

40.3.4.1 Software Inventory Management. Enterprise patch management is dependent on having a current and complete inventory of the patchable software (applications and operating systems) installed on each host. This inventory should include not only which software is currently installed on each host, but also what version of each piece of software is installed. Without this information, the correct patches cannot be identified, acquired, and installed. This inventory information is also necessary for identifying older versions of installed software so that they can be brought up to date. A major benefit of updating older versions is that it reduces the number of software versions that need to be patched and have their patches tested.

40.3.4.2 Resource Overload. Enterprise patch management can cause resources to become overloaded. For example, many hosts might start downloading the same large patch (or bundle of patches) at the same time. This could consume excessive network bandwidth or, if the patches are coming from an organization patch server, overwhelm the resources of that server. Organizations should ensure that their enterprise patch management can avoid resource overload situations, such as by sizing the solution to meet expected volumes of requests, and staggering the delivery of patches so that the enterprise patch management system does not try to transfer patches to too many hosts at the same time.

40.3.4.3 Installation Side Effects. Installing a patch may cause side effects to occur. A common example is the installation inadvertently altering existing security configuration settings or adding new settings. This may create a new security problem in the process of fixing the original vulnerability via patching. Organizations should be capable of detecting side effects, such as changes to security configuration settings, caused by patch installation.

40.3.4.4 Patch Implementation Verification. As discussed in Section 40.3.1 of this chapter, an installed patch might not take effect until the affected software is restarted or other state changes are made. It can be surprisingly difficult to examine a host and determine whether or not a particular patch has taken effect. This is further complicated when there is no indication for a patch when it would take effect (reboot required/not required, etc.). One option is to attempt to exploit the vulnerability, but this is generally only feasible if an exploit already exists, and there are substantial risks with attempting exploitation, even under highly controlled conditions. Organizations should use other methods of confirming installation, such as a vulnerability scanner that is independent from the patch management system.

40.3.4.5 Application Whitelisting. Application whitelisting technologies can conflict with patch management technologies because the application whitelisting technologies function based on known characteristics of executables and other application components, which may be changed by patching. If the vendor is providing the whitelist information, the vendor will have to acquire the patch, record its files' characteristics, and send the corresponding information to customers. If the organization

is building its own whitelist information, it will have to acquire each patch, record its files' characteristics, and update its whitelists with the new information. Either method may cause problematic delays for organizations that apply patches quickly, especially automatically; patched software may be seen as unknown software and prohibited from running.

To avoid these problems with updates, most application whitelisting technologies offer maintenance options. For example, many technologies allow the administrator to select certain services (e.g., patch management software) to be trusted updaters. This means that any files that they add to or modify on a host are automatically added to the whitelist. Similar options are available for designating trusted publishers (i.e., software vendors), users (such as system administrators), sources (such as trusted network paths), and other trusted entities that may update whitelists. Organizations using application whitelisting technologies should ensure that they are configured to avoid problems with updates.

40.4 ENTERPRISE PATCH MANAGEMENT TECHNOLOGIES.

This section provides an overview of enterprise patch management technologies. It discusses their composition, focuses on the security and management capabilities that they provide, and gives recommendations for their use.

40.4.1 Components and Architecture.

Enterprise patch management technologies are similar architecturally to other enterprise security solutions: one or more centralized servers that provide management and reporting, and one or more consoles. Enterprise patch management technologies can also be offered as a managed service. What distinguishes enterprise patch management technologies from each other architecturally are the techniques they use to identify missing patches. The three prevalent techniques are agent-based, agentless scanning, and passive network monitoring. Many products support only one of these techniques, while other products support more than one. All the techniques are explained in more detail in the following subsections. Organizations should carefully consider the advantages and disadvantages of each technique when selecting enterprise patch management technologies.

40.4.1.1 Agent-Based.

An agent-based patch management technology requires an agent to be running on each host to be patched, with one or more servers that manage the patching process and coordinate with the agents. (Note that agent-based patch management technology is built into some operating systems.) Each agent is responsible for determining what vulnerable software is installed on the host, communicating with the patch management servers, determining what new patches are available for the host, installing those patches, and executing any state changes needed to make the patches take effect (e.g., application restart, OS reboot). Each agent runs with administrator privileges so it can perform these actions. The patch management server is responsible for providing the agents with information on vulnerable software and available patches, including where patches can be acquired and what state changes are needed.

Compared to agentless scanning and passive network monitoring, agent-based patch management technologies are strongly preferred for hosts that are not on the local network all the time, such as telecommuter laptops and smartphones.

There are a few limitations to agent-based patch management technologies. Hosts that don't permit direct administrator access to the operating system, such as many

appliances, generally cannot run agents. Also, agents may not be available for all of the organization's platforms.

40.4.1.2 Agentless Scanning. An agentless scanning patch management technology has one or more servers that perform network scanning of each host to be patched and determine what patches each host needs. Generally, agentless scanning requires the servers to have administrative privileges on each host, so that they can return more accurate scanning results and they have the ability to install patches and implement state changes on the hosts (application restarts, OS reboots, etc.).

The main advantage of agentless scanning is that it doesn't require the installation and execution of an agent on each host.

One of the primary limitations of agentless scanning is that it omits hosts not on the local network, such as telecommuter laptops and mobile devices. Also, network security controls (e.g., host-based firewalls) and network technologies (e.g., network address translation) may inadvertently block scanning or otherwise negatively affect scanning results. Agentless scanning may also negatively impact operations by consuming excessive amounts of bandwidth. Finally, agentless scanning may not support all of the organization's platforms.

40.4.1.3 Passive Network Monitoring. Passive network monitoring technologies for patch management monitor local network traffic to identify applications (and in some cases, operating systems) that are in need of patching.

These technologies can be effective at identifying hosts that are not being maintained by other patch management solutions (agent-based, agentless scanning). They do not require any privileges on the hosts to be monitored, so they can be used to monitor the patch status of hosts that the organization does not control (unmanaged systems, visitor systems, contractor systems, etc.).

The primary disadvantage of passive network monitoring is that it only works with software where you can identify the version based on its network traffic (assumed to be unencrypted). Also, of course, it only works with hosts on the local network.

40.4.2 Security Capabilities. This section describes common security capabilities provided by patch management technologies, divided into three categories: inventory management, patch management, and other.

40.4.2.1 Inventory Management Capabilities. Patch management technologies typically have capabilities for identifying which software and versions of software are installed on each host, or alternately, just identifying vulnerable versions of software that are installed. In addition, some products have features for installing new versions of software, installing or uninstalling software features, and uninstalling software.

40.4.2.2 Patch Management Capabilities. Patch management technologies obviously provide a range of patch management capabilities. Common features include identifying which patches are needed, bundling and sequencing patches for distribution, allowing administrators to select which patches may or may not be deployed, and installing patches and verifying installation. Many patch management technologies also allow patches to be stored centrally (within the organization) or downloaded as needed from external sources.

40.4.2.3 Other Capabilities. Many host-based products that have patch management capabilities also provide a variety of other security capabilities, such as antivirus software, configuration management, and vulnerability scanning. Further discussion of these capabilities is outside the scope of this chapter.

40.4.3 Management Capabilities. Once a patch management technology has been selected, its administrators should design a solution architecture, perform testing, deploy and secure the solution, and maintain its operations and security. This section highlights issues of particular interest with the management—the implementation, operation, and maintenance—of patch management technologies, and provides recommendations for performing them effectively and efficiently.

40.4.3.1 Technology Security. Deploying enterprise patch management tools within an enterprise can create additional security risks for an organization; however, a much greater risk is faced by organizations that do not effectively patch their systems. Such tools usually increase security far more than they decrease security, especially when the tools contain built-in security measures to protect against security risks and threats. The following are some risks with using these tools:

- A patch may have been altered (inadvertently or intentionally).
- Credentials may be misused.
- Vulnerabilities in the solution components (including agents) may be exploited.
- An entity could monitor tool communications to identify vulnerabilities (particularly when the host is on an external network).

Organizations should reduce these risks through the application of standard security techniques that should be used when deploying any enterprise-wide application. Examples of countermeasures include the following:

- Keeping the patching solution components tightly secured (including patching them)
- Encrypting network communications
- Verifying integrity of patches before installing them
- Testing patches before deployment (to identify corruption)

40.4.3.2 Phased Deployment. Organizations should deploy enterprise patch management tools using a phased approach. This allows process and user communication issues to be addressed with a small group before deploying the patch application universally. Most organizations deploy patch management tools first to standardized desktop systems and single-platform server farms of similarly configured servers. Once this has been accomplished, organizations should address the more difficult issue of integrating multiplatform environments, nonstandard desktop systems, legacy computers, and computers with unusual configurations. Manual methods may need to be used for operating systems and applications not supported by automated patching tools, as well as some computers with unusual configurations; examples include embedded systems, industrial control systems, medical devices, and experimental systems. For such computers, there should be a written and implemented procedure for the manual patching process.

40.4.3.3 Usability and Availability. Organizations should balance their security needs with their needs for usability and availability. For example, installing a patch may "break" other applications; this can best be addressed by testing patches before deployment. Another example is that forcing application restarts, OS reboots, and other host state changes is disruptive and could cause loss of data or services. Again, organizations need to balance the need to get patches applied with the need to support operations. A final example, particularly important for mobile devices, is the acquisition of updates over low-bandwidth or metered connections; it may be technically or financially infeasible to download large patches over such connections. Organizations should make provisions for ensuring that their enterprise patching solution works for mobile hosts and other hosts used on low-bandwidth or metered networks.

40.5 METRICS AND MEASURES. As explained in Section 3.3 of NIST SP 800-55 Revision 1, Performance Measurement Guide for Information Security, there are three types of measures:

1. Implementation measures are used to demonstrate progress in implementing security programs, specific security controls, and associated policies and procedures…
2. Effectiveness/efficiency measures are used to monitor if program-level processes and system-level security controls are implemented correctly, operating as intended, and meeting the desired outcome…
3. Impact measures are used to articulate the impact of information security on an organization's mission…[4]

Regarding these types of measures, "less mature information security programs need to develop their goals and objectives before being able to implement effective measurement. More mature programs use implementation measures to evaluate performance, while the most mature programs use effectiveness/efficiency and business impact measures to determine the effect of their information security processes and procedures." Accordingly, organizations should implement and use appropriate measures for their enterprise patch management technologies and processes.

Examples of possible implementation measures include:

- What percentage of the organization's desktops and laptops are being covered by the enterprise patch management technologies?
- What percentage of the organization's servers have their applications automatically inventoried by the enterprise patch management technologies?

Examples of possible effectiveness/efficiency measures include:

- How often are hosts checked for missing updates?
- How often are asset inventories for host applications updated?
- What is the minimum/average/maximum time to apply patches to X% of hosts?
- What percentage of the organization's desktops and laptops are patched within X days of patch release? Y days? Z days? (where X, Y, and Z are different values, such as 10, 20, and 30).

- On average, what percentage of hosts are fully patched at any given time? Percentage of high impact hosts? Moderate impact? Low impact?
- What percentage of patches are applied fully automatically, versus partially automatically, versus manually?

Examples of possible impact measures include:

- What cost savings has the organization achieved through its patch management processes?
- What percentage of the agency's information system budget is devoted to patch management?

40.6 FURTHER READING

Jang, M. (2006). *Linux® Patch Management: Keeping Linux® Systems Up To Date*. Prentice-Hall, 2006.

Nicastro, F. M. *Security Patch Management*, 2nd ed. CRC Press, 2013.

Taylor, J. R., J. H. Allen, G. L. Hyatt, and G. H. Kim. *Change and Patch Management Controls: Critical for Organizational Success*. Global Technology Audit Guide: The Institute of Internal Auditors, Inc., 2005. www.aicpa.org/interestareas/informationtechnology/resources/itassuranceservices/guidance/downloadabledocuments/gtag2patchmgmtcontrols.pdf

USGAO. *Information Security: Continued Action Needed to Improve Software Patch Management.* U.S. Government Accountability Office, 2011.

USGAO. *Information Security: Effective Patch Management is Critical to Mitigating Software Vulnerabilities.* U.S. Government Accountability Office, 2011.

40.7 NOTES

1. This chapter is based on the National Institute of Standards and Technology (NIST) Special Publication 800-40 Version 2.0, *Creating a Patch and Vulnerability Management Program*, and Special Publication 800-40 Revision 3, *Guide to Enterprise Patch Management Technologies*, which are available at http://csrc.nist.gov/publications/nistpubs/
2. NIST SP 800-53 Revision 4 is available for download at http://csrc.nist.gov/publications/PubsSPs.html#800-53-rev4
3. The PCI DSS is available for download at www.pcisecuritystandards.org/security_standards/
4. NIST SP 800-55 Revision 1 is available for download at http://csrc.nist.gov/publications/PubsSPs.html#800-55

CHAPTER 41

ANTIVIRUS TECHNOLOGY

Chey Cobb and Allysa Myers

41.1 INTRODUCTION **41·1**
- 41.1.1 Antivirus Terminology 41·2
- 41.1.2 Antivirus Issues 41·3

41.2 ANTIVIRUS BASICS **41·3**
- 41.2.1 Early Days of AV Scanners 41·4
- 41.2.2 Validity of Scanners 41·4
- 41.2.3 Scanner Internals 41·5
- 41.2.4 Antivirus Engines and Antivirus Databases 41·6

41.3 SCANNING METHODOLOGIES **41·7**
- 41.3.1 Specific Detection 41·7
- 41.3.2 Generic Detection 41·8
- 41.3.3 Heuristics 41·8
- 41.3.4 Intrusion Detection and Prevention 41·9
- 41.3.5 Reputation-Based Scanning 41·9

41.4 CONTENT FILTERING **41·10**
- 41.4.1 How Content Filters Work 41·11
- 41.4.2 Efficiency and Efficacy 41·11

41.5 ANTIVIRUS DEPLOYMENT **41·12**
- 41.5.1 Desktops Alone 41·12
- 41.5.2 Server-Based Antivirus 41·12
- 41.5.3 Mobile Devices 41·13

41.6 POLICIES AND STRATEGIES **41·13**

41.7 CONCLUDING REMARKS **41·13**

41.8 FURTHER READING **41·14**

41.1 INTRODUCTION. For over three decades, malware has been a persistent, annoying, and costly threat, and there is no end in sight to the problem. There are many vendors offering to provide a cure for viruses and malware, but the mere existence of these software pests is understandably vexing to those charged with system security.

Initially, most viruses were not designed to cause harm but were created more to gain notoriety for the creator or as a prank. Because these early viruses were designed to subvert legitimate program operations across multiple systems, they were more likely to cause unexpected problems because the virus writers didn't do exhaustive testing. These viruses, and later some Trojans, often damaged data and caused system downtime. The cleanup required to recover from even a minor virus infection was expensive in terms of lost productivity and unbudgeted labor costs.

Viruses and Trojan behavior have merged, and now both are considered as part of the larger family referred to as malware. No longer is malware just written for a virus writer's 15 minutes of fame; today, malware is created primarily for financial gain. Malware can still cause damage, but now it is more likely to have been created

to steal valuable information or the resources of an infected computer. Malware can steal valuable data for sale on underground markets, send spam from a compromised machine, or install software designed to eavesdrop on an affected user's activities, to name a few examples. Some malware programs can give full access to a hacker to control infected machines, and malware continues the trend of controlling affected machines in a "botnet." These collections of hijacked computers are used for a variety of activities, such as launching attack or sending spam. A botnet makes use of the increased power of networked computers, while making it more difficult to trace a suspected attacker. For more information about malware, see Chapters 16, 17, and 20 in this *Handbook*.

Because of this shift to a financial motivation for malware writing, something of an open-source community has sprung up to develop new malware. There have been two major areas where development has been most notable: evasion of traditional antivirus defenses and distribution by exploiting vulnerabilities in common software. In their efforts to evade traditional antivirus defenses, criminals often slightly modify and compress malware with *packers*, which frequently use sophisticated obfuscation or anti–reverse-engineering techniques. The primary focus of today's malware writers is to avoid the newsworthy outbreaks of the past, staying under the radar to accomplish targeted attacks (*phishing* and *spear phishing*). Some of the most common malware families now have as many as tens of thousands of variants with hundreds of variants being released each day. Once they are in a system, they often employ stealth techniques such as kernel-mode rootkits to hide themselves from the operating system and from many antivirus (AV) products.

New malware programs are being developed every day because vulnerabilities in software programs are so numerous, and it is often weeks between the time a vulnerability is discovered and the time that the vendor issues a security patch. Even more often, patches are announced by vendors but not implemented by users. Malware writers take advantage of these vulnerabilities and time lags, so research into vulnerabilities is imperative in order to create effective defenses with antivirus programs. A layered defense against malware infection is essential and should employ a modern antivirus scanner as well as a firewall and data encryption.

Although the threat of new malware remains, and its growth has become exponential, the technology deployed to defend computers against viruses and malware is one of the least understood aspects of security architecture. As a result, antivirus defenses are often set up improperly. This chapter describes available antivirus technologies and how they work and outlines methods for their effective application within a comprehensive program of computer security.

41.1.1 Antivirus Terminology. The acronym *AV* is widely used to describe the industry, the products (also sometimes abbreviated *AVPs*), and the programs that have been developed to defeat computer viruses.

Early AV programs used simple hard- and floppy-drive scans to search for a specific text string hidden within a specific file or a boot sector. This is the origin of the term *AV scanner*, which is now widely used as a generic term for all AV programs. That is how the term is used in this chapter, although it should be pointed out that many of today's AV scanners do much more than merely scan for known malware, and detection may now include unwanted programs as well as suspicious file activities.

When exploring AV product literature, it is common to see statements about the number of viruses and Trojans that exist. It is not unusual to see estimates of the total number of malware variants in the hundreds of millions, and six-digit estimates

of the number of new malware added every day. AV vendors' methods for counting differ from one to the next, and these numbers tell us very little about their detection capability. It is far more instructive to rely on the reports of reputable AV testing labs to get a clear idea of a product's effectiveness, than to rely on these numbers.

41.1.2 Antivirus Issues. In the early days of viruses, aside from having the ability to spread to other machines, there was little benefit in keeping a virus on a machine for any length of time. During the early days, viruses were relatively easy to find, because a virus usually caused noticeable damage that made users aware that their systems were infected. At other times, an unmistakable message or image appeared on the screen, alerting users to the infection. As virus writers discovered a monetary benefit to remaining on a system as long as possible, the number of affected machines increased tremendously. Viruses and malware have become complex, are adept at disguising themselves, and are not as easy to find and eradicate. To counter this problem, AV scanners have become sophisticated logic machines. This new level of complexity, of both the problem and the solution, has made it even more confusing for users. Too many people fail to understand the importance of having up-to-date AV products. Users do not know what to look for, and do not understand that viruses and malware do not announce their presence.

Modern AV software has an automatic-update capability, which has made updating a less labor-intensive process. Updates are the key to keeping systems malware-free because updates give the AV program the information necessary to find the newest viruses and malware. Unfortunately, because malware is released at such a prodigious pace, even daily updates of a traditional string-based scanner may not be enough to keep systems totally clear of threats. Many AV products now support essentially constant contact with the developers' servers and verify checksums on the servers to identify changes requiring a download. These methods allow updates within minutes of any change with no inconvenience to the users.

Most AV scanners look at what has happened before (i.e., known behavior of known malware), in order to detect infections. Although AV products now attempt to anticipate new threats, that is not their strong point. Additionally, many consumers who buy a new computer receive a working AV product with their purchase, but they do not choose to buy a subscription for AV updates once their free use period expires. They do not understand why they have to pay for something that seems to be working, and they are lulled into a false sense of security by thinking that their AV product is doing its job. Consumers who bring work from unprotected home computers into the office can quickly infect an entire network.

Security products historically have suffered from a lack of upper-management support because they are often viewed as high-cost/low-return items, and they are given a low priority in the security budget. The information technology (IT) team of an organization needs resources for things like identifying and patching vulnerable systems, monitoring and managing machines entering their network, monitoring suspicious behavior identified by firewalls, AV products, and security scanners. Considering the wide array of AV products from which to choose, and the large number of patches introduced each day, many system administrators despair. They will install and patch just what they have time for, and consider it good enough.

41.2 ANTIVIRUS BASICS.

AV scanners are a bit like police officers walking the beat. They try to watch everything that is going on around them, look out for suspicious behavior, and attempt to intercede when they think something bad is happening or about

to happen. Both the police and AV scanners look for certain patterns and behaviors, and they leap into action when a suspect crosses a predetermined threshold of acceptability. Like the police, AV scanners sometimes reach the wrong conclusions. These errors are usually caused by insufficient data or by new and unexpected behavioral patterns.

Malware detection is an inexact science, and it is impossible to create an AV scanner with a 100 percent success rate. It is simply not possible to know the intent of every bit of code that enters a computer, and it is not feasible to test every bit of code before it executes. To do so would require that the AV scanner demand so much of the processing power of the CPU that valid programs would not be able to execute. Malware behaviors are subject to broad variations, and many use stealth or cloaking techniques to hide from the operating system and even from the AV scanner itself. There are no longer hard-and-fast rules that a user can apply to determine if a system harbors malware.

41.2.1 Early Days of AV Scanners. When viruses first started appearing with regularity in the late 1980s, their detection and eradication was relatively straightforward, but not necessarily easy. The viruses were quite simple and normally did not spread very quickly. The AV community quickly researched them, determined what made them work, and published effective fixes in short order. These fixes tended to be written for a specific virus and could not disinfect other viruses, even if they were of similar types. Most of the work fell on users to identify which virus (or type of virus) they thought they had and then search for a program that would fix it. Because Internet connectivity was not as prevalent as it is now, users frequently spent much time calling friends and associates in the hope that they could forward a disk copy of the necessary AV program. Additionally, there were no naming conventions for viruses, and it was difficult to determine with any conviction that the fix obtained would actually work.

Viruses of that period generally inserted their code in predictable sections of a program. The early scanners ran a search for a specific string of characters. If they found it, they would delete the virus code and attempt to restore the host program to its original uninfected form. Failing that, the scanner would usually advise the user that disinfection was incomplete and that they should delete the infected application and reinstall it.

As the number of viruses began to climb, software companies that had ventured into the AV market began to realize that creating and distributing individual fixes was no longer feasible. Instead, they began to develop more comprehensive scanners that could look for more viruses, both old and new. The new generations of scanners were comprised of two components: the scanning engine and the signature files. Each component was entirely dependent on the other to work. The engine consisted of the user interface and the application that scanned the system for viruses. The signature files were a database of the *fingerprints* (unique segments of code) of known viruses. Although some of these early scanners did a good job, many did not. None of the early AV scanners was able to catch all known viruses.

41.2.2 Validity of Scanners. The vendors of software scanners in the late 1980s and early 1990s faced a number of obstacles. It seemed there was a new AV vendor appearing every month, and the market became highly competitive as user awareness of the virus problem grew. Given this competitive state, there was vast dissension among the AV community as to how viruses for research should be stored and tested. Many AV vendors kept a library of viruses for their own use, and this fact was used in their marketing. Claims that one program worked better than another because it checked for more viruses were misleading because no one knew how many

viruses existed. There was simply no method of commercial or independent testing to check the validity of claims made by the AV product vendors. Additionally, there was a problem of naming the viruses. Each vendor created its own names for viruses, and it was not uncommon for one virus to be known by several names.

The AV vendors also disagreed on how AV scanners should operate in principle. Some vendors felt that AV scanners should only look for new viruses, and others felt that a good product should search for both old and new viruses. While this argument raged, viruses looked as though they would eventually gain the upper hand, especially as virus writers began to use underground bulletin boards, and later the Internet, to share and distribute virus code.

With no standards for the AV products, the public had little to go by other than the vendors' marketing copy and the advice of other users. However, if a recommendation was made by a friend for Brand X Antivirus because no viruses were found on the friend's system, it was possible that no viruses had ever been introduced in the friend's system at all, and Brand X could not find old viruses, new viruses, or any viruses at all.

Two things happened that revolutionized the AV scanner market. In 1993, Joe Wells, a research editor with a business magazine, began collecting viruses and virus reports from experts around the world and began assembling a library of these viruses. He named this library of viruses the *WildList* and made it available to legitimate AV researchers. His list divided viruses into those known to have infected systems (in the wild) and those that had been written but were not actively infecting (in the zoo). A naming convention of viruses also began to emerge in order to maintain an efficient and searchable database.

The other notable event was the development of commercial AV testing and certification by a company known as the National Computer Security Association (NCSA), which is now known as ICSA Labs. The NCSA started a consortium of AV vendors that, for a fee, submitted their products to be tested. NCSA and Joe Wells began collaboration for the use of his WildList, and Dr. Richard Ford, a noted virus expert, created a virus-testing laboratory for NCSA. Dr. Ford fashioned an environment in which AV scanners were put through their paces to see if they could detect all of the viruses in the WildList. AV vendors submitted their products every time a new version of their product was about to be released. Although the original test results were dismal (many scanners could not detect more than 80 percent of the viruses in the list), an environment had been created in which measurable improvements in the effectiveness of AV technology could be achieved. Naturally, the public and the press began to look for AV products that had been certified by NCSA. Eventually, other commercial and independent test laboratories independently developed their own certification and testing schemes to help users find reliable AV products.

Before long, other testing organizations joined the fray and AV products increased in complexity. It was clear that there was a need for testing standards to help users distinguish effective, unbiased tests from those that did not clearly demonstrate the capabilities of AV products. A group of AV researchers, academics, testers, and other interested parties created an organization to develop such standards, called the Anti-Malware Testing Standards Organization (AMTSO). This group has published a number of papers outlining considerations for how to produce or distinguish a good test from an inadequate one. These papers are meant to help both testers and those people reading tests.

41.2.3 Scanner Internals. As was noted earlier, an AV scanner cannot simply put each program into a computer's RAM and test it for malware before the program is

allowed to execute. To do that would require almost all the resources of the CPU, and users would have a system that operated at a snail's pace.

In order to operate efficiently, AV scanners have had to go to great lengths to improve their speed in order to check for the vast number of existing malware without bringing the entire system to a halt. Using sophisticated elimination methods to exclude signatures whose characteristics do not match those of the file being scanned, helps keeps speeds reasonable.

There are five basic types of detection within a modern AV product. Most people are familiar with signature-based detection, which is also known as specific detection, but this is not the only technique in place anymore. AV products may use some or all of these five basic methods of operation:

1. **Specific detection**. Looking for known malware
2. **Generic detection**. Looking for infections by variants of known malware
3. **Heuristics**. Scanning for previously unknown viruses by noting suspicious behavior or file structures, as described more fully in Section 41.3.3
4. **Intrusion prevention**. Monitoring known-suspicious system changes and behaviors to prevent suspected malware
5. **Reputation**. Detecting malware by a file or Website's reputation, based on a number of factors including the number of times the file or Website has been downloaded, customer reports, or age of a Website

41.2.4 Antivirus Engines and Antivirus Databases. The AV engine and its signature database work in concert to prevent and detect malware trying to enter a system. The engine generally provides a library of commonly used functions. It consists of dozens of complex searching algorithms, CPU emulators, and various forms of programming logic. The engine determines which files to scan, which functions to run, and how to react when suspected malware is found. However, the engine knows absolutely nothing about the malware themselves and is almost useless without the signature database.

The signature database contains the fingerprints (snippets of distinctive code) of hundreds of millions of malware variants. As new malware and variants appear at an accelerating rate, it is imperative that the signature database be updated often. In 1995, the experts advised updating the database files at least once a month, but with so many viruses appearing each day, users today are advised to update at least daily—and ideally, to allow constant updates controlled by their AV product. AV manufacturers now provide products that check for updates automatically and download changes whenever a user is connected to the Internet. Some vendors also provide cloud-based signature databases so that computers that are connected to the Internet can check against a more-frequently updated set of signatures that are hosted on vendors' own remote systems.

The signature database also contains the rule sets used in heuristic scans. These types of scans can be slower and more intrusive than simple signature scans, and their design and implementation vary greatly between products. Most products now give users configurable options to lessen or increase heuristics as desired. Although signature scans can be considered a heuristic in themselves, the term is more commonly used to identify the more complex AV functions that attempt to locate viruses by identifying suspicious behavior and/or file structure.

Because the distinction between a scanning engine and a signature database is not obvious to many system administrators, many religiously update the database but are unaware that the engine also may need updating. This is a poor strategy that can result in many viruses slipping by the scanner undetected.

41.3 SCANNING METHODOLOGIES. AV products are configurable by the user or the system administrator to scan upon startup, constantly, or on demand. To be at its most effective, a scanner should be set to a continuous or "on access" scan, with a periodic, scheduled scan set to occur when the system is on but not in use. Users running less-optimized AV programs may find that this degrades system performance. Some scanners need to use much of the system's memory on continuous scans in order to be able to test sections of code, which may make the applications noticeably slower when they first run. Therefore, a happy medium must be found—the AV scanner must be able to protect the system, while the user must be able to have full use of the system.

There is no one scanning method that is superior to the others. All of the scanning methods have their advantages and disadvantages, but none is able to detect malware with unfailing accuracy. A scan is looking for code and behaviors that have been noticed in other malware, and if a new malware exhibits new, previously unknown behaviors, it can pass by undetected. Therefore, most AV scanners do not rely on only one scanning method to detect malware, but have several included in their design.

41.3.1 Specific Detection. Each malware uses different code to perform its functions. A sequence of code that is specific to each malware is referred to as the fingerprint, or signature of that malware. To detect the presence of malware, the scanner looks for the signature, removes its code from the host file or system, and attempts to restore the infected program or system to an uninfected state. In early viruses, it was discovered that the signatures were usually found within specific areas of a program, specific to each virus. The scanners set out to inspect only those areas of a file rather than scanning an entire program from top to bottom. This saved vast amounts of time and processing power.

As malware is often static, rather than parasitic in nature, AV researchers now must get more creative to quickly identify known-bad files. But as malware authors have armored their files to make their creations difficult for researchers to analyze, this can also make a malicious file look very different from a valid document or application. As such, researchers can key in on these differences and swiftly exclude benign files.

Every vendor's AV product has a different implementation of scanner and database, although the signature scanning technique is the most common. Signature scans can identify whether a program contains one of the many signatures contained in the database, but it cannot say for certain whether a system has actually been affected by malware (e.g., the malware may be present but not yet executed). Users can only trust the guess of the AV scanner, because the odds are in the scanner's favor. It is possible, however, that a program that is suspected of being infected actually contains random data that only coincidentally looks like a virus signature. The legitimate program could contain instructions that by sheer chance matched the search string in the virus database. However, when there is a possibility that the code is actually from a virus, the scanner reports it as a positive hit.

False-positive reports are a possible problem of signature scanners. If users notice that their scanner falsely reports the presence of viruses too often, they view this as an annoyance, and will likely seek to disable the software or find ways of circumventing the scans.

41.3.2 Generic Detection. As was discussed earlier, malware is now often created for financial purposes, and it makes fiscal sense for its authors to get as much use as possible out of successful creations. They often make open-source code that is shared widely in the malware-author community, so that it is often updated with new functionality such as exploit-code or password-stealing capabilities, to target new games, applications, or online banking sites.

Likewise, it makes sense for AV scanners to look for common properties of popular malware families or known malicious behavior in order to proactively detect variants based on those codebases. These generic detections can vary from being more heuristic in nature to being fairly specific. For example, it could include protection as broad as buffer overflow detection to prevent certain types of exploits, to specific homegrown packers used by only one malware family.

Generic detection caused a great deal of controversy in the early days of AV products. There was concern, when most viruses spread by infecting clean files parasitically, that generic detection would require generic cleaning, which could leave host files mangled beyond repair or usability. As the vast majority of malware now infects systems rather than host files, simple deletion of malicious files and cleaning of registry entries can remove malware. But because malware is also frequently component-based, removing one threat may be only the tip of the iceberg. There may be other components that are as yet undetected, and it's best to thoroughly examine any system which has had malware detected. As such, more and more system administrators have adopted a policy of simply re-imaging affected systems rather than relying on AV products' removal procedures.

41.3.3 Heuristics. By adding heuristics to their AV scanners, vendors looked to increase the efficacy of their products. The scanners could now look for malware that are new and unknown and not contained within the signature database.

The word *heuristic* comes from a Greek word meaning "to discover." The term is used today in computer science to describe algorithms that are effective in solving complex questions quickly. A heuristic algorithm makes certain assumptions about the problem it is trying to solve. In the case of an AV scanner, it analyzes a program's structure, its attributes, and its behavior to see if these meet the rules that have been established for identifying malware, even without its specific signature being known. Basically, a heuristic algorithm works on the assumption that if it "looks like a duck, walks like a duck, and sounds like a duck, it must be a duck."

The drawback to heuristic scanning is that it makes intelligent assumptions but is nevertheless bound to make mistakes. Another problem with heuristic scanning is that, on slower systems, it may take longer to run and may require user interaction. Some users consider this intrusive and may turn off the feature. By combining both signature scanning and heuristic scanning in their products, AV vendors have increased their effectiveness and speed.

Heuristic scanners use a rule-based system to verify the existence of malware. It applies all the rules to a given program and gives the program an overall score. If the score is high, there is a good likelihood that it is malicious. Generally, the scanner first looks for features of a file that are more common to malicious files. A well-designed heuristic scanner will limit the regions of the program to be examined in order to scan the highest number of suspects in the shortest possible time. The scanner then examines the logic of the suspected program to determine if it might be trying to perform known, likely malicious actions. This type of scanning is considered to be a *static* scan. The

static method applies the rules and gives a pass/fail score to the program—whether the program has actually executed or not.

The other type of heuristic scanning is called the *dynamic method*. This method applies basically the same rules as the static method, and if the score is high, it attempts to emulate the program. Rather than examining the logic of the suspected code, the dynamic scanner runs a simulation of the virus in a virtual environment. This technique has come to be known as *sandbox* emulation, and is effective for attempting to identify new malware that does not appear in the signature database.

Neither of these heuristic scanning methods is necessarily better than the other, but in concert they give fairly good results. Although a static heuristic scan may miss some malware because they have not yet executed, the dynamic heuristic scan can catch previously unknown malware before they are allowed to execute on the system. Between heuristic and generic detections, it is becoming increasingly common for at least one AV vendor to identify any new malware as soon as it is released.

41.3.4 Intrusion Detection and Prevention. As more complex malware appears at an increasingly prodigious pace, scanning programs solely for signatures has become less effective at finding viruses. Virus authors use encryption or obfuscation techniques, or release a large number of individual variants in the hopes that the AV scanner will not find them. Today's operating systems and legitimate programs have bloated to millions of lines of code, so that finding a virus signature may be resource intensive. Because malware may steal valuable data or damage systems, it is not a good strategy to let them execute and then attempt to clean up the mess. A better strategy is to try to find malware before it has had a chance to affect a system and to prevent it from doing harm.

The use of a cyclical redundancy check (CRC), or checksums, was originally added to some security products, including AV suites, to aid in the detection and prevention of virus infections. This method tracks unauthorized file and system changes, as when malware or a hacker enters and alters a system. To track those changes, a fingerprint of each executable program and data file is computed and stored in a database when the AV product is first installed. These fingerprints are quite small, usually consisting of less than 100 bytes of information—this is the "sum" or checksum. Because malware must add or change files on a system in order to affect it, the checksums of the fingerprints are compared with any newer version. If the checksums vary, then the AV scanner runs other routines to investigate further.

In the case of intrusion-prevention systems, a behavior-based scan is executed on each new file when it is run. If certain suspicious behaviors are observed, a user may be asked whether the behavior should be allowed to continue. If a sequence of behaviors is observed that is sufficiently malicious, the program may be halted entirely. This technique has been found to be remarkably effective in preventing execution of new, unknown malware, although it shares the same sort of difficulties found with using heuristic scans. Again, as part of a complete security arsenal, it can be a valuable tool. For more information on intrusion detection and intrusion prevention, see Chapter 27 in this *Handbook*.

41.3.5 Reputation-Based Scanning. As new malware is released at such a prodigious rate, it is no longer practical for researchers to investigate each sample individually. While tens or even hundreds of thousands of samples are investigated daily, there are still a number of threats that slip through the cracks. In order to deal with this constant deluge, rather than relying solely on a handful of researchers to

look for known-malicious files or Websites, companies are beginning to rely on the general public to increase their view of the malware landscape. Reputation scanning gathers telemetry and reputation data provided by a product's users, plus information about the age of files and Websites, and then rates them as trusted, bad, or somewhere in-between. This can help a user determine whether to proceed, based on the relative risk level.

This sort of scanning has limitations. Certain types of files are unique by design, especially data files such as documents, spreadsheets, and presentations. And if a user chooses to run a file that turns out to be malicious, they're on their own when it comes to figuring out how to remove the threat from their system. But this is why AV products include a variety of different scanning technologies; to provide protections that will ideally overlap in such a way that they can catch a larger percentage of brand-new, malicious code.

41.4 CONTENT FILTERING. In the early days of virus infections, computer security experts often allayed the fears of computer users by telling them that they could never catch a computer virus from email. This assurance was based on the fact that email was almost exclusively composed of ASCII text documents, with no ability to execute program code. At the same time, the skeptics were saying "never say never." The skeptics won.

First, there were several waves of macro virus–infected documents sent as file attachments to email. This led to the modified assurance from security experts that no one could ever catch a computer virus from an email message attachment, if the attachment was not opened. Then virus writers started embedding commands that use HTML and the scripting capability of email programs. This led to the further modified assurance from experts that a computer virus could not be caught from unopened email. This assurance in turn proved unwarranted, because email preview capabilities were exploited to trigger malicious code even without user intervention. At one point, merely highlighting a message subject in MS Outlook was enough to execute an attachment, although this default was later changed.

Virus writers also began to exploit the user's email facility by forwarding copies of the virus to entries in the user's email address book. The Melissa virus was the first virus that really leveraged email to spread rapidly. Since the Melissa virus, users have been advised to suspect just about any unsolicited email. Malware writers are always looking for new delivery methods, and they were richly rewarded when email programs began to allow executable code within the email. Although users enjoy the point-and-click convenience of this feature, it allowed new viruses to proliferate at a rate not seen before.

The Web also has seen an explosion in malicious code distribution, particularly since the advent of "Web 2.0"— social networking and collaboration sites. Through Adobe Flash, Java, and JavaScript objects that exploit vulnerabilities in Internet browsers or media-viewer software, malware can be automatically downloaded and installed on users' machine without their being aware. This is commonly referred to as a *drive-by download*.

Content filtering is one popular way of controlling Web and email threats. It consists of a server-based application that interrogates all incoming and outgoing traffic, according to its configuration and rule sets. Early versions were cumbersome to configure, due to a text-based interface that required all rules to be composed in a laborious text editor. Misconfigurations were commonplace because administrators often were not sure which of the text files was causing a failure.

The new generation of content filters has increased user friendliness, using interactive graphic user interfaces to set and adjust policies. Administrators are able to fine-tune policies so that they meet the specific needs of their organizations. For example, all email containing executable attachments may be blocked, quarantined, or simply deleted. This may be established as a rule, for some users or all users.

Content filters were particularly effective in preventing infections of email-borne viruses, even before the specific virus signature was released. For example, when the security officer of one government office heard of the Love Bug virus on the morning news, that person set the content filters to block all email attachments containing the extension ".vbs," thus averting infection of a large network. No user interaction was required, and most users were not aware that the block had been placed. The costs and downtime of a virus attack were prevented.

The majority of threats now come from compromised legitimate Websites rather than the dark, seedy underbelly of the Internet, which changes the way products view unwanted content. Consequently, much of this content filtering technology has moved from a focus on blocking potentially malicious content to blocking Websites that are deemed inappropriate for viewing at work or by children. Many content filters come with extensive, frequently updated block-lists for porn, "hate-speech," and other politically sensitive subjects.

41.4.1 How Content Filters Work. These applications work in the same general manner that AV scanners do. They scan all incoming data on specific ports on the server, and they compare the traffic to rules and strings in the database. Because content filters are capable of blocking more than one type of file or program, they have the ability to scan text files, graphics, zipped files, self-extractors, and various executables. Many content filters contain AV scanning components, so if traffic does contain known-malicious code, it can be intercepted and disinfected before it is sent to the recipient.

The standards for formatting email for transmission using standard protocols have long been in place and detail the type of information in every section. It is easy for a program to look for particular information within these sections to determine what is included in the message—attachments, for example. Content filters first begin by disassembling a message to look at its various parts before scanning the message for the items to be allowed or denied into the system. Before sending the message onward, it is reassembled and checked for the conditions specified in the configurations. For example, a condition may state that any attachments be stripped and deleted, that the message body be sent to the recipient, and that an outward email be sent to the sender stating that attachments are not allowed in email.

In terms of data security, a content filter adds several elements that are beyond the traditional AV scanner. For example, message and attachment content can be scanned for inappropriate material. This might be proprietary company information that employees should not be sending out via email, or it could be offensive material such as pornography, that should not be coming into the company's email system. Content filtering also can stop the spread of common forms of spam mail, such as chain letters and get-rich-quick schemes.

41.4.2 Efficiency and Efficacy. Speed of operations is a concern with content-filtering mechanisms, given the large volume of network traffic in most organizations. However, because the operations are all contained within the server, the users will not notice any change in performance of their desktop systems; that is, the processing will

be completed before the messages are received by the client systems. For instance, in the case of email, if filtering causes mail to queue up, delivery of mail may lag for a period of time. If Websites are blocked for inappropriate content, the traffic will simply not be delivered to the user.

Content filters are also subject to the same failures as traditional AV scanners. New malware can be missed if the data is not present in the scanning database. Additionally, the configuration of the product and the application of patches and updates are crucial to its successful operation. False positives are also a problem, where a legitimate message inadvertently includes content that triggers a block. It is possible to quarantine questionable messages and have a system administrator follow up with the sender. This can lead to refinement of the filters or to the detection of serious offenses. Before a content-filtering system is deployed, it is important to put in place the response mechanisms, so that abuses of policy can be addressed appropriately. This may well involve several departments besides security, including legal and human resources.

41.5 ANTIVIRUS DEPLOYMENT. AV scanners can be installed on the desktop or on servers. Each strategy has its advantages and disadvantages. For example, if the system is server based, malware on USB thumb drives, DVDs, and CDs on the desktop will not be scanned. The consensus of most experts, however, is to use both. With the advances in AV products and network management systems, it is entirely possible to install scanners both on the desktop and on servers while still maintaining an acceptable level of control and performance.

41.5.1 Desktops Alone. If an organization's computer security policy allows unrestricted use of thumb drives, floppies, DVDs, and CDs, it imperative that AV scanners be deployed to the desktop. Unless these drives are locked or disabled, there is no way, other than scanning, to prevent users from accidentally introducing malware. The preferences of a desktop AV scanner can be set to automatically scan external media.

Updates to desktop AV scanners can now be distributed via a central server. This is particularly effective when new signature files are needed to prevent infiltration by newly discovered malware. The updates can be pushed to the desktop, and the users need not be present at the workstation, although the desktop system must be on and connected to the network at the time. If the updates are scheduled after working hours, it is important to verify that systems have updates pushed to them as soon as they log back into the network. Some AV products have management consoles that will allow this to occur automatically, rather than having someone check each system individually.

In order to prevent unauthorized AV-setup changes, it is possible to prevent the users from changing the configuration of their desktop AV scanners. Since that may not be the default installation, it must be checked. Again, with the use of a management console, it is possible to enforce security-software policies remotely.

41.5.2 Server-Based Antivirus. Many companies have sought to bolster the defenses at the perimeter of their network by installing AV products on the server where downloads are frequently stored and traffic is high. A server-based AV scanner can be configured to send alerts to administrators when suspected malware is detected. Like the desktop-based scanners, the response to malware detection can be predetermined. Many system administrators set the program to erase all infected files rather than to send them to quarantine. This strategy works to lessen the possibility that quarantined malware can be "released" by mistake.

41.5.3 Mobile Devices. With the increase in use of smartphones and tablets by individuals, organizations, and through bring-your-own-device policies, securing these devices against malware has grown in importance. Host-based intrusion detection and signature-based tools have been developed to protect such devices, but there are problems based on the security model of some operating systems (e.g., Android) because of the strict partition of memory areas controlled by applications.

41.6 POLICIES AND STRATEGIES.

In the battle against malware, the promulgation of appropriate policies and the implementation of realistic plans of action are equally as important as the installation of an AV scanner. The policies should spell out in detail what actions are allowed or denied, and they should also be specific about the users' responsibilities. The policies and responsibilities should be regularly updated to reflect the realities of the changing malware landscape, especially within a company's ecosystem.

End-user AV awareness training should be high on the list of priorities in every organization. Users are more likely to cooperate in preventing and quarantining infections if they are aware of the types of malware that may affect their system and of the damage they can cause. A simple security-incident bulletin board in a central location is an easy and effective way to communicate with users. Email is probably not an effective method of distributing malware awareness information because users become confused between the education effort, actual and legitimate malware alerts, and bogus malware alerts.

The roles and responsibilities of each person within an organization should be clearly defined and communicated to the general populace. For example, the responsibilities of an average user will be different from those of a system administrator, and the responsibilities should be reflected in their roles. An individual user's role may describe the actions required of the user if malware is detected on a workstation, while the system administrator's role may describe how to handle the report from the user and prepare for removal.

Problems and catastrophes will usually occur when they are least expected, so every organization should have an emergency response plan in its policies. The emergency plan should detail the list of persons to be called in an emergency and the priority order in which they should be called.

For every major malware incident within an organization, care must be taken that a "lessons learned" session be undertaken as soon after the event as possible. No matter how well-written a policy may be, it cannot be proven effective until it is put into use. An actual infection will highlight the failures of a policy in action, which should be rectified before the next attack. For more on security policy guidelines see Chapter 4 in this *Handbook*.

For more details of computer security incident response team management, see Chapter 56 in this *Handbook*.

Management's support for such policies is vital. Support is required not only to approve the AV budget and the policies, but also to ensure that everyone abides by the policies. It is highly unlikely that users will follow a policy that upper management routinely flouts.

41.7 CONCLUDING REMARKS.

Both AV technology and malware technology have progressed rapidly over time, and AV technology largely remains in step with malware technology. The success of malware is now primarily due to financial

motivation, but a large part of the reason it is such a lucrative business is that people continue to dismiss the threat of malware. However good AV technology gets, it will not make a serious dent in the malware problem unless it is appropriately implemented by organizations and properly employed by users who act responsibly. As AV technology continues to improve and becomes better understood, we hope that it will continue to be used more widely and more wisely.

41.8 FURTHER READING

AMTSO. "Documents and Principles." Anti-Malware Testing Standards Organization, 2013. www.amtso.org/documents.html

AVG. www.avg.com/us-en/homepage

Avast. www.avast.com/index

Aycock, J. D. *Computer Viruses and Malware*. New York: Springer, 2006.

Bitdefender. www.bitdefender.com

Drake, J. J., Z. Lanier, C. Mulliner, P. Oliva, S. A. Ridley, and G. Wicherski. *Android Hacker's Handbook.* Wiley, 2013.

Elisan, M. *Malware, Rootkits & Botnets A Beginner's Guide*. McGraw-Hill Osborne Media, 2012.

ESET. www.eset.com/us

Johnston, J. R. *Technological Turf Wars: A Case Study of the Antivirus Industry*. Philadelphia: Temple University Press, 2008.

Kaspersky Lab. http://usa.kaspersky.com

Li, Q., and G. Clark. " Mobile Security: A Look Ahead." *IEEE Security and Privacy* 11, no. 1 (January 2013): 78–81.

Ligh, M., S. Adair, B. Hartstein, and M. Richard. *Malware Analyst's Cookbook and DVD: Tools and Techniques for Fighting Malicious Code*. Wiley, 2010.

Malin, C. H., E. Casey, and J. M. Aquilina. *Malware Forensics Field Guide for Windows Systems: Digital Forensics Field Guides*. Syngress, 2012.

McAfee. www.mcafee.com/us

Miller, C., D. Blazakis, D. DaiZovi, S. Esser, V. Iozzo, and R.-P. Weinmann. *iOS Hacker's Handbook.* Wiley, 2012.

Nazario, J. *Defense and Detection Strategies against Internet Worms*. Artech House, 2003.

Scott, W. *Virus and Malware Removal Made Easy.* Amazon Digital Services, 2013.

Sikorski, M., and A. Honig. *Practical Malware Analysis: The Hands-On Guide to Dissecting Malicious Software*. No Starch Press, 2012.

Skoudis, E., and L. Zeltser. *Malware: Fighting Malicious Code*. Prentice Hall, 2003.

Slade, R. "Antiviral Software Evaluation FAQ." Doug's Home on the Web, 2006. www.dmuth.org/virus/faqs/evaluation.html

Symantec | Norton. http://us.norton.com

Szor, P. *The Art of Computer Virus Research and Defense*. Addison-Wesley, 2005.

Vasudevan, A., J. M. McCune, and J. Newsome. *Trustworthy Execution on Mobile Devices.* Springer, 2013.

Virus Bulletin. www.virusbtn.com/index

CHAPTER 42

PROTECTING DIGITAL RIGHTS: TECHNICAL APPROACHES

Robert Guess, Jennifer Hadley, Steven Lovaas, and Diane E. Levine

42.1 INTRODUCTION 42·1
- 42.1.1 Digital Rights 42·2
- 42.1.2 Patent, Copyright, and Trademark Laws 42·2
- 42.1.3 Piracy 42·2
- 42.1.4 Privacy 42·3

42.2 SOFTWARE-BASED ANTIPIRACY TECHNIQUES 42·3
- 42.2.1 Organizational Policy 42·4
- 42.2.2 Software Usage Counters 42·4

42.3 HARDWARE-BASED ANTIPIRACY TECHNIQUES 42·5
- 42.3.1 Dongles 42·5
- 42.3.2 Specialized Readers 42·6
- 42.3.3 Evanescent Media 42·10
- 42.3.4 Software Keys 42·11

42.4 DIGITAL RIGHTS MANAGEMENT 42·12
- 42.4.1 Purpose 42·12
- 42.4.2 Application 42·13
- 42.4.3 Examples 42·13

42.5 PRIVACY-ENHANCING TECHNOLOGIES 42·14
- 42.5.1 Network Proxy 42·14
- 42.5.2 Hidden Operating Systems 42·14

42.6 POLITICAL AND TECHNICAL OPPOSITION TO DRM 42·15
- 42.6.1 Political Opposition 42·15
- 42.6.2 Technical Countermeasures 42·17

42.7 FUNDAMENTAL PROBLEMS 42·17

42.8 SUMMARY 42·18

42.9 GLOSSARY 42·19

42.10 FURTHER READING 42·22

42.11 NOTES 42·22

42.1 INTRODUCTION. Ever since publishing and commerce were introduced to the digital world, the risks to intellectual property and to personal privacy in cyberspace have steadily escalated on comparable but separate paths. These paths have now converged. Unfortunately, many times, antipiracy efforts lead to possible breaches in personal privacy.

Efforts to stem the flow of pirated software worldwide remain mediocre in efficacy; piracy is still proving to be big business in the new millennium. According to the Business Software Alliance (BSA), the "global piracy rate hovered at 42 percent in 2011 while a steadily expanding marketplace in the developing world drove the commercial

value of software theft to $63.4 billion." Further, "these startling findings come from a survey of approximately 15,000 computer users in 33 countries that together make up 82 percent of the global PC market."[1] Continuing piracy means lost jobs, wages, tax revenues, and a potential barrier to success for software start-ups around the globe.

At the same time, freedoms inherent in the Internet have made maintaining privacy of personal information a true challenge. Identity theft keeps rising, as more and more companies actively engage in the accumulation of customer data through e-commerce. By paying bills, checking medical insurance accounts, or completing taxes online, people are making their personal information available to be stockpiled, shared, and regurgitated to the point that simply "Googling" one's identity can be a real eye opener. Technologies that aim to prevent piracy have the potential to use this wealth of available personal information in a way that significantly erodes personal privacy. In particular, some of these technologies send personally identifiable information to servers on a routine basis (see, e.g., Section 42.4.2). The idea that a corporation—or a government—could be scanning what a specific person reads, listens to, or views is grounds for concern to civil libertarians.

42.1.1 Digital Rights. In this environment of rapid change in both piracy and privacy, even the term *digital rights* is ambiguous. When software companies and music producers talk about digital rights, they mean the kinds of rights long protected by copyright, trademark, and patent law. When privacy advocates argue about digital rights, however, they may be talking about a completely different thing: that an individual does not forfeit personal rights, including the right to privacy, merely by turning on a computer.

This chapter's primary focus is on technologies designed to protect traditional rights of content producers, but it also enumerates areas where those technologies threaten personal privacy.

42.1.2 Patent, Copyright, and Trademark Laws. There are differences among the applicable laws and the materials they protect.

Patents give owners exclusive rights to use and license their ideas and materials; patents generally protect nonobvious inventions in mechanical and electrical fields, as well as those that can be embodied in computer software and hardware.

Copyrights give owners the exclusive rights to create derivative works, to reproduce original works, and to display, distribute, and conduct their works. Copyrights apply to original works of authorship including paintings, photographs, drawings, writings, music, videos, computer software, and any other works that are fixed in a tangible medium. Copyrights, their infringement, and remedies are described in the Copyright Act of 1976.

Trademarks give owners the right to restrict the use of distinctive marks in certain contexts. These rights may apply to words, sounds, distinctive colors, symbols, and designs.

For an extensive discussion of intellectual property law, see Chapter 11 in this *Handbook*.

42.1.3 Piracy. Once thought of as a mere copyright infringement of printed matter or production of a counterfeit audiotape, piracy has grown with technology and has expanded to encompass intellectual property, digital data, DVDs, CDs, VHS, analog and high-definition TV, and streaming media.

There are several types of piracy. End user piracy occurs when end users use a single copy of software to run on several different systems, or when they distribute copies of software to others without permission of the software manufacturer. Reseller piracy occurs when unscrupulous resellers distribute multiple copies of a single software package to multiple customers, preload the same software on multiple systems, or knowingly sell counterfeit software to customers. Internet and bulletin board (BBS) piracy occurs when users download and upload copyrighted materials and use it or make it available for use by others without proper licenses.

To understand why and how piracy occurs and the enormous impact on society worldwide, we need to have a clear understanding of what we mean by the word "piracy." Whenever information is created and published in print, on the Internet, or incorporated into software, that information may be protected by copyright, patent, or trademark law. This principle applies to a broad spectrum of material that includes, for example, the Wright Brothers' specifications for their "Flying Machine"; Microsoft Windows software; the icon Mickey Mouse and all related materials; and television shows, plays, movies, and music created and performed live and on recordings. Making unauthorized copies of such material in any medium is referred to as *piracy*.

Within the last three years, the Software & Information Industry Association[2] has "filed more than 100 lawsuits in the U.S. against illegal eBay sellers as well as sellers on other websites dealing in counterfeit, OEM, academic, region-specific, and other illegal software and publications. Defendants have paid millions of dollars in damages, and, in some cases, criminal charges were pursued and defendants sentenced to jail time."[3]

42.1.4 Privacy. As in many aspects of security, end users are now being called on to protect their online (and offline) identities through diligent monitoring of financial activities and through awareness programs focused on personal rights, laws, and obligations of Internet use. The average Web user today can create and publish a blog or personal video to the Web faster than that same user can apply software patches to a desktop PC. So widespread is this new facility that the law has yet to catch up with the needs and expectations of the users in online society. It is difficult to keep data private when—with a mouse click—it is easy to share personal information with the whole world.

Beyond bad personal habits that reduce privacy, however, a whole new class of applications termed *digital rights management* (DRM) collects more personal information than ever before in an effort to reduce improper use of copyrighted material. DRM products may record and report on an individual's Web-browsing habits, types of files created and accessed by a particular program, number of uses of a particular file or program, source IP address of the user's system, and presence (or absence) of a license for a program. In the name of protecting digital rights for content producers, the consumers of this content are being cataloged and tracked in a way that the framers of copyright laws would have been hard-pressed to imagine.

For extensive coverage of privacy issues on the Internet, see Chapter 69 in this *Handbook*.

For a glossary of terminology used in discussing digital rights management, see Section 42.9.

42.2 SOFTWARE-BASED ANTIPIRACY TECHNIQUES. A variety of software-based technical approaches are used to prevent inappropriate use of copyrighted and otherwise protected material in an organization's networks and on the

public Internet. Current methods include ensuring proper configuration of operating systems, monitoring of installed software, encryption of content, and insertion of some sort of key or identifier in the digital product itself.

42.2.1 Organizational Policy. Controls in free-standing commercial applications represent only one of the technical means to protect digital content. Existing system controls are also useful in this regard. Operating system access controls can specify who can access particular content, while encryption supported by the OS and other applications can limit access to users who possess the appropriate key. More generally, organizational policy should specify good practices in configuring operating systems and applications in order to help protect content. Such policy includes:

- Allow users to install only software that is necessary.
- Encrypt information that should not be publicly viewable.
- Install software with the lowest possible privilege consistent with the ability to do its job.
- Disable active content (Java, JavaScript, ActiveX, cookies, etc.) wherever feasible.
- Use network operating system access controls to limit access to shared, copyrighted media to members of the organization for whom licenses are purchased.

42.2.2 Software Usage Counters. Software metering has been popular for several years. Special software monitors system usage and inventories the software on the system or network. This type of software also can be used to block or limit the use of specific software, such as browsers and games. In addition to fighting piracy, it can reduce the load on IT personnel by reducing complications due to use of unauthorized software.

42.2.2.1 Controlling Concurrent Installations. Software metering products can monitor concurrent installations even on networks used by people in different geographic areas who have different requirements and different software installed. The metering software permits an administrator to maintain a live and updated inventory of the software installed at different locations on the network. The logs show where installation has taken place, when the licenses expire, and when updates are necessary. Alerts can be set to notify system administrators when a license is about to expire or when an update has been accomplished.

42.2.2.2 Controlling Concurrent Usage. Software metering allows network administrators to identify and resolve cases of illegal installation of unauthorized copies of authorized software and also to catch people who install unauthorized software on the organization's computers. Metering software also allows a company to report and analyze logon and logout times, track software usage, and meter software licenses to keep those in the company legal. In addition to avoiding legal entanglements, monitoring can reduce the demand on system resources, network bandwidth, and technical support staff.

42.2.2.3 Examples and Implementation. Microsoft announced in 2000 that to combat piracy, new releases of the Office 2000 program would include a counting feature making the programs malfunction if the owner had not registered the software after launching it 50 times.[4] Prior to this announcement, Microsoft published

antipiracy literature and provided a great deal of consumer education regarding the effects of software piracy on society; it established, as well, a piracy hotline (1-800-RU-LEGIT). Novell (1-800-PIRATES), Adobe, and Xerox are other companies that have vigorous antipiracy programs in place, although they have not yet announced that they are building metering into their products.

Metering software requires a bona fide license in order for it to be implemented. Typically, companies establish CD-ROM keys that are printed on legitimate copies of their product installation discs or jewel cases. The keys include checksums or message authentication codes that can be checked by the installation routines. The algorithms for the checksums are intended to cause difficulty for people trying to create counterfeit keys. The security of such measures depends on the cryptographic strength of the validation keys.

Software counters for controlling concurrent usage need to store information securely so that each load operation increments a counter and each unload operation decrements it. However, the security problem is to store this information in such a way that unauthorized people cannot modify it easily. Encrypting the data in a complex sequence of operations can stop most abuses of the system by making the effort required for circumventing the mechanisms more costly than buying a license.

More recently, starting with Vista, copies of the Windows operating system must be registered (Microsoft calls this *activating* a copy) soon after installation, or they will go into *reduced functionality mode*, in which basic functions are available for only an hour of operation before logging on again.[5] For single-copy home use, Windows relies on a license key encoded on the installation medium. For enterprise deployments of Vista, an organization needs to run key servers to allow the registration to occur.[6]

42.3 HARDWARE-BASED ANTIPIRACY TECHNIQUES.

Working on the theory that software is prone to misconfiguration and compromise, antipiracy groups and researchers have experimented with a variety of hardware-based approaches to preventing inappropriate use of protected content. These techniques include dongles and specialized readers attached to the reading hardware, evanescent media designed for viewing or playing only a limited number of times, and software keys incorporated into the media and the reading hardware.

42.3.1 Dongles.

Dongles are hardware lock devices or modules that connect to a computer and communicate with software running on the computer. Without a dongle in place, the external device or regulated software does not work fully, or at all.

Initially, dongles controlled printing. With a dongle installed on the computer, no one could print data from the computer without authorization. However, there is now a necessity for protecting all types of devices. Now dongles are used to protect scanners; external drives (e.g., ZIP drives); CD-ROMs and rewritable CD-ROMs; DVDs and DVD-Rs; VHS recorders; PlayStation, Nintendo, and Sega video gaming systems; and even personal digital assistants (PDAs).

The most common type of dongle provides a pass-through port to connect a device cable. Generally, a dongle incorporates some type of algorithmic encryption in its onboard microelectronic circuitry. The sophistication of the encryption varies depending on the manufacturer and the device. Many dongles provide additional onboard nonvolatile memory for the software to access. Some models even have real-time checks that keep track of date and time information, including when an application's license (temporary or leased) is set to expire.

Dongles provide some definite advantages:

- Because a dongle is an external device, it is fairly simple to install and uninstall. Early dongles used serial or parallel ports, but USB has become the norm in recent years. In most cases, since manufacturers support their devices, an ordinary user can install and use a dongle without help from an IT department.
- Dongles also require registration, which provides adequate control over the use of the dongle and thus provides legitimacy to both the device and the users. Registration (dependent on the contract in place) may provide support for both the software and the hardware.
- Dongles that support encryption provide an extra layer of protection by making transmitted data indecipherable until it reaches its destination, unless the hardware is in place.

There are also disadvantages to using dongles:

- Consumers resist the requirement for installation, maintenance, and additional cost. Most large corporations do not use dongles for their products.
- Dongles can be lost or stolen, and they also may fail.
- Sometimes a dongle will work well with a slow computer but cause errors when installed on a faster computer.
- Since not every manufacturer automatically replaces lost or stolen dongles without charge, there may be additional costs involved in getting replacements.
- Dongles can present a serious risk-management problem for critical applications where delays in obtaining replacements or registering them may be unacceptable.
- As with any device, there can be a serious problem if the dongle manufacturer ceases to support the model of dongle a company has installed or if the manufacturer goes out of business entirely.
- Laws regarding encryption usage differ in various countries. Specialized dongles that may be legal to use in the United States may be illegal in another country.

42.3.2 Specialized Readers. One of the impediments to illegal copying used to be the difficulty and cost of obtaining specialized hardware and software for reading and copying proprietary materials with fidelity. However, today such copying equipment is inexpensive and easy to find. In addition, the media for distributing illegal copies are less expensive than ever.

42.3.2.1 Audio. According to the Recording Industry Association of America (RIAA), the global audio industry loses in excess of $4 billion every year to piracy worldwide.[7] RIAA says that $1 million a day, in just physical product, is lost in the United States alone. The loss of ancillary revenues drives the figures higher. But the RIAA claims that these figures are low, since it estimates that in some countries up to 98 percent of the music in use comes from illegal copies.

As part of an industry-wide, organized approach to highlighting and reducing the music piracy problem, the RIAA has taken a very active role in pursuing legal action against suspected pirates. Through the 1990s and continuing into the current decade, the biggest problem with audio piracy was illegally copied CDs. In 1998, for example, the RIAA confiscated 23,858 illegal CDs in the first half of the year. In that same year,

Operation Copycat—a joint investigation by RIAA, the Motion Picture Association of America (MPAA), and the New York Police Department—saw the arrest of 43 CD pirates and the shutdown of 15 illegal manufacturing locations. Many of the CDs seized in these types of operations apparently came from Asia and Eastern Europe, financed by organized crime operations with ties to drugs and prostitution.[8] By 2002, even corner convenience stores were contributing to the problem, providing coin-operated CD copying machines akin to photocopiers, along with the familiar posted warning transferring liability to the user.[9] Given the enormous profits involved, the problem has been simply too large and widespread for law enforcement to control. In 2005, the RIAA seized approximately 5 million illegal CDs.[10]

More recently, the RIAA has focused on the problem of music files illegally shared across the Internet. Using free software, sometimes already bundled into commercial operating systems, anyone can download music tracks and burn CDs. The MP3 music file format has become a ubiquitous way to share music with others. The RIAA originally protested against MP3 players, but the phenomenal success of personal digital music players, particularly the Apple iPod, has rendered such efforts futile. Some musicians and independent record labels have adopted the MP3 format to promote their records, showing that the technology itself has no inherent ties to piracy. These musicians and recording studios claim that they are happy with consumers downloading the music, and they are at odds with the RIAA.

Some musical groups have even experimented with severing ties with the traditional music industry altogether, using Websites and social networking services to advertise and distribute digital music files directly to their listeners.[11] As more musicians begin to use the Internet either in addition to or in lieu of traditional distribution models, the very model of music distribution is in flux.

Meanwhile, the music industry continues to face serious financial loss due to illegal downloading. The industry, through the RIAA, has been pursuing legal remedies. Some major lawsuits have achieved significant press coverage and have been instrumental in enforcing or changing existing laws or in helping develop new laws. For instance, deliberate copyright violations resulted in an award of $50 million in statutory damages and $3.4 million in legal fees to Universal Studios in a suit against the MyMP3.com music service. MyMP3 created a database of over 80,000 albums, which, when combined with the MyMP3 software, let users access and store music digitally, without paying a fee.

Perhaps the most recognizable name in the music field in regard to piracy at one time was Napster, a site that enabled individuals to share tracks of music via the Internet. The site provided free downloadable software for downloading and playing MP3 files. Essentially, Napster software turned a user's PC into part of a distributed server network that published available music files. It did not take long for the site to acquire a user group of millions of people who after "sampling" the music might then go to a store and buy the entire CD.

However, since the Napster site did not limit the length of the download, many users simply downloaded the entire track. Most never bought the commercial version of the music. Adding to successful piracy attempts was the development and ready availability of rewritable CD drives. More and more Napster users decided to download the music tracks they wanted and then burn their own CDs without ever purchasing the CDs made by the recording artists and music companies.

Creating a stir in the industry and ultimately a landmark judicial case, Napster was forced to radically alter operations in March 2000 after protracted court proceedings. Upon the verdict, Jack Valenti, president and chief executive officer of MPAA,

commented that the consumer would benefit most from the court's decision because "You cannot take for free what belongs to someone else."[12] But the subject of Napster and audio piracy remains highly controversial. Although some people argued that little-known artists received exposure that they may never have gotten without the free file-sharing service, others, especially large music companies and recording artists, argue that they were being denied the royalties they deserve.

Napster attempted to re-create itself as a pay-for-subscription music download service, but found record labels unwilling to work with it. In 2002, Napster folded; the name was eventually purchased by Roxio, Inc. to rebrand its own subscription service. In the meantime, other peer-to-peer (P2P) file-sharing protocols and applications appeared to fill the void left by Napster. By the time the iTunes Music Store emerged as a powerful contender with music company backing and with copyright protection, the world of free file sharing was reinvigorated by names such as Gnutella, FastTrack, Grokster, Limewire, and Kazaa.

Over the last few years, the music industry has identified universities as fertile ground for pursuing illegal downloading activities. In 2007, the RIAA launched a new round of attempts to bring music pirates to justice, sending letters offering to settle with students identified as probably sharing copyrighted files, in advance of any trial.[13] The RIAA's tactics have been raising hackles in the higher education community, with opponents criticizing the letters as bordering on extortion. Some universities are refusing to forward the letters to students; one university agreed to forward the letters, but promised to bill the RIAA $11 for every letter to pay for its staff time.[14]

42.3.2.2 Video. On the video side, Scour, Inc., provided free downloads of digital movies as well as software allowing users to share the downloaded files among themselves without the use of a central server. With an easy-to-use interface and quick response time, Scour.com became quite popular in a short period of time. Launched in 1997, with a Web search feature added in 1998 and a subsequent P2P tool, Scour eventually attracted negative attention from the movie and music industries. In July 2000, MPAA, RIAA, and the National Music Publishers Association (NMPA) sued Scour, accusing it of large-scale theft of copyrighted material and of trafficking in stolen works. By November 2000, the company was out of business.[15]

This case did not stop the counterfeiting of video media. Despite increasingly steep fines, judicial rulings, and even raids by various law enforcement agencies, counterfeit video is readily available. Along Fifth Avenue in New York City, for $5 to $10 anyone can buy the latest films and DVDs; in markets in Hong Kong, Southeast Asia, and India, the copies are even cheaper. It is true that some of the copies available may have been "legally produced," but it is more than likely that the counterfeit or bootleg copies were made illegally from a master copy that was either borrowed or stolen.

Advances in consumer electronics help the trend, as many PCs now come with a recordable/rewritable DVD player as a standard component. Arguments about legality of time-shifting and space-shifting that once defended the practice of making personal mixes on audiocassette have now moved to the realm of digital video.

42.3.2.3 Television (Analog). Broadcast television has been one of the most successful technologies in history, and financial interests are still huge, despite the growth of cable and satellite services. In January 2000, major television companies, the National Football League, and the National Basketball League all filed complaints against iCraveTV, a Canadian company that had been in existence for only a year.

According to the complaints, iCraveTV was illegally using broadcast television signals without authorization or payment and streaming the signals to the iCrave Internet site for viewing free of charge. Although this practice apparently did not violate Canadian copyright laws at the time, U.S. judges ruled in February 2000 that the unauthorized transmissions of broadcast signals into the United States via the Internet were a direct violation of U.S. copyright law, and iCraveTV was ordered to stop the practice. Shortly after iCraveTV agreed to an out-of-court settlement, the Web-site was shut down and iCraveTV went out of business.[16]

Hacking cable decoders is another technique for obtaining services without paying for them. Although it is not illegal to buy, install, or modify equipment for converting encoded cable TV signals from pay-per-view or other commercial suppliers, it is illegal to use such set-top decoders to obtain services without paying for them.

In the United States, Congress mandated that after February 17, 2008, all TV stations must transmit in digital format (DTV) only.

42.3.2.4 Television (HDTV). The first television image was created in 1884 when Paul Nipkow created a mechanical scanning disk. With only 18 lines of resolution, the picture was poor. Current National Television System Committee (NTSC) standard TV transmissions are done with bandwidth that does not exceed 6 MHz. The current analog system broadcasts 30 frames per second and 525 lines per frame.

High-definition television (HDTV) is a digital television system that offers twice the horizontal and vertical resolution of the current TV system. HDTV has the ability to deliver a video composed of approximately 1,125 lines per frames and 60 frames per second. Viewers then see a picture quality close to that of 35-mm film. Obviously, transmitting images containing that large amount of audio and video information requires wide bandwidth, actually about 18 MHz. Such bandwidth would permit the transmission of 1,050 lines of 600 pixels per line. However, the Federal Communications Commission (FCC) decided to limit HDTV to a 6-MHz maximum bandwidth. In order to meet that requirement, MPEG compression would be used.

MPEG compression applies algorithms to pixel groups and records information that changes within a frame, rather than all of the information in all of the frames. Audio is synchronized to the video. Using MPEG saves storage space and transmission requirements while retaining high image and sound quality. According to the Advanced Television Committee Standard (ASTC), the FCC will require that audio and video compression as well as the transmission of HDTV terrestrial signals follow this standard.

As with all other transmissions and media, there are serious concerns about piracy of HDTV transmissions and programs. At the present time, even though many TV transmissions are scrambled in order to thwart reception, it is fairly simple, although illegal, to purchase a descrambler and unscramble the transmissions. It is, however, legal for home viewers to record programs for their own personal use.

The HDTV market space has been evolving, and consumer demand for HD-capable devices has exploded over the past several years. Attempts by U.S. television producers to protect themselves and their content using a variety of scrambling and encryption schemes, including content scrambling systems (CSSs), have been made difficult by the frequent changes in hardware and signal formatting that have accompanied this rapid market expansion.

Encrypting terrestrial broadcast television programming would secure the transmissions, but according to the Home Recording Rights Coalition (HRRC), such encryption will threaten established home recording rights. The HRRC contends that Section 1202

(k) of the Digital Millennium Copyright Act provides a carefully balanced approach to analog home recording rights and stipulates that mandated technology may not be applied to interfere with consumer recording of free, over-the-air terrestrial broadcasts.

Furthermore, the HRRC contends that encrypting the free television broadcast content will create very little incentive for consumers to switch from regular analog to digital television. Instead of thwarting digital pirates, the HRRC contends that strong encryption will impose unfair and even illegal restrictions on consumers.

42.3.2.5 Consumer Acceptance of Specialized Readers. Illegal sharing of copyrighted content is common across the Internet. When the Software Publishing Association (SPA) was first formed, the group, together with law enforcement agencies, raided the physical premises of companies believed to be using pirated software. As a result of finding quantities of the pirated software, SPA won many legal actions and related settlements.

Some people see encryption as a challenge and work at breaking the algorithms so they can pirate the data—digital, video, or audio. In addition, the lack of standardization of laws throughout industries and countries has led to controversy and ongoing piracy.

Although average consumers do not think of themselves as intellectual property pirates, many otherwise honest citizens do get and use illegal programs, applications, games, audio tracks, CDs, DVDs, VHS tapes, and television signals. This situation might be attributed to a lack of ethical education, but many people like to save money and simply do not believe they will be caught and punished for such pilfering. A 2001 study by the Pew Internet & American Life Project, based on phone interviews with 4,205 adults 18 and over, some 2,299 of whom were Internet users, suggested that around 30 million U.S. residents had downloaded music from the Internet.[17] At that time, the phrase "had downloaded music from the Internet" was basically equivalent to "had *illegally* downloaded music from the Internet," since legal means of doing so had yet to evolve.

Since that report, Apple's 2001 announcement of its iTunes and iPod products, and the 2003 launch of the iTunes Music Store, began a movement to provide consumer-friendly ways of downloading music that also give compensation to copyright owners. iTunes uses device authentication and proprietary encoding formats to limit redistribution of downloaded songs and videos. Despite—or perhaps because of—Apple's attempts to comply with copyright laws, it is clear that Apple filled a perceived need in the market, as it has generated both a host of competitors and substantial sales. Consumers downloaded the first million songs from the iTunes Store in five days, and the overall market for digital music has grown to at least $790 million per year. At $0.99 per song, downloads from the iTunes Music Store passed the $1 billion mark on January 23, 2006.[18]

42.3.3 Evanescent Media. There are many interpretations of the term *evanescent media*. The broad interpretation includes digital imaging, optics, multimedia and other electronic art, and data that are short-lived or transitory. When such media are original, creative works, society has an interest in protecting them against piracy.

Since most evanescent media involve some visual aspects as well as text, antipiracy techniques now being used or considered for other types of data may be applicable. Such techniques include previously discussed dongles, software keys, watermarks, encryption, and digital rights management. Part of the problem in electing and implementing a solution is the lack of existing standards that specifically deal with this new area of art and science.

42.3.4 Software Keys. Software keys of various kinds are used to secure data and equipment. A software key is generally a string of numbers that is used for identification purposes, either to allow access to the use of equipment or to permit authorized printing, processing, or copying of data. As described earlier in the discussion of dongles, most anticopying hardware devices are accompanied by software that works in tandem with the hardware. A software key activates or deactivates the hardware lock. When the software is working perfectly, there are generally no difficulties. However, all software can malfunction, and when that happens, there can be serious problems in getting equipment to work. Additional problems occur when the computer containing the software key malfunctions, and the software key cannot be made to work on a replacement machine.

42.3.4.1 Videocassettes versus Copy Machines. Watermarking is one of the techniques being seriously considered for protecting videocassettes and DVDs. In 1995, the ASTC formed a Copyright Protection Technical Working Group, which spun off a special Watermarking and Embedded Data Encoding subgroup. The group has broad representation, including representatives from the PC market, the Macintosh market, the MPAA, the Consumer Electronics Manufacturers Association (CEMA) and related manufacturers, technicians, and users. Their task is to look for technologies and services that might use hidden data clues as a means of inhibiting or barring digital piracy. Using a hidden watermark that can be embedded in the content would then prevent machines from making copies or would alert the operator that the videocassette is marked and that unauthorized copies would be considered pirated.

42.3.4.2 DVD Area Encoding. Digital video requires very large storage space—too large for a single CD to hold. However, by applying compression techniques, the digital video can be compressed to fit into the digital videodisc's maximum capacity of 17 gigabytes. Two different types of compression are used for encoding audio and video content for DVD: constant bit rate (CBR) and variable bit rate (VBR) compression.

In order to prevent piracy of the content of the DVD, many companies are turning to encryption. The compressed data are encapsulated through a mathematical algorithm that can be decrypted only through the use of a decryption key.

42.3.4.3 Implementation. For shorter programs, CBR is ideal. Based on MPEG-2 encoding, CBR compresses each frame of audio and video by a user-selected amount. This degree of compression is then applied to the entire program. Using VBR, it is possible to create a database of video content based on the amount of change in each frame or scene. This is particularly useful in programs with a long format. To construct the database, the encoding software does several analytical passes of the master footage and then makes a final digitizing pass. From the created database, the computer can encode the video with a variable data rate, allowing a higher bit rate for scenes with pans, zooms, and fast motion and giving scenes with little or no motion low data rates. By greatly compressing the areas of lower detail, areas of higher details can be allocated more space and use less compression.

42.3.4.4 Watermarks. Watermarking involves embedding one set of data inside a larger set of data. The embedded set of data identifies the origins or ownership of a specific work, just as a watermark does on paper.

Using digital watermarks can help copyright owners track the use of anything digital, including music, movies, photographs, and clip art. Digital watermarking is widely used for protecting images. For instance, photographers often post low-resolution (low-res) versions of their photos on public Websites and use visible digital watermarks to clearly label the low-res images as copyrighted. Upon payment of the appropriate fee, the customer receives the high-resolution version of the photo, presumably with at least any visible watermarks removed. The use of invisible watermarks to prevent undetected sharing after purchase of digital content is more controversial and more prone to questions about reliability of detection; for instance, how many false positives and false negatives will occur? Additionally, there is the question of survivability of the mark itself as it is run through various transformations.

The music industry flirted with digital watermarking to protect music files beginning in 1998 with the formation of the Secure Digital Media Initiative (SDMI), a consortium of technology, security, and music organizations. The SDMI developed several watermarking schemes, and in 2000, it offered a reward to anyone who could crack the code and remove the watermark from a song protected by SDMI's technologies. The Electronic Frontier Foundation asked the Internet community to boycott the contest, stressing that the use of DMAT (Digital Music Access Technology) would mean that manufacturers and users would be forced to adopt the DMAT format in equipment and would create additional costs for manufacturers and consumers. A team of researchers, led by Princeton professor Ed Felten, was able to remove the invisible watermarks. When Felten attempted to publish the results of his process, attorneys for SDMI threatened to sue him under the Digital Millennium Copyright Act (DMCA). SDMI never filed suit, but Felten himself sued for a declaratory judgment to clarify the matter. Felten's suit was dismissed by a federal judge, but not before the government and the RIAA agreed that researchers should not be punished under the DMCA for testing technologies to protect copyright.[19] The SDMI has been inactive since 2001.[20]

42.4 DIGITAL RIGHTS MANAGEMENT. Recognizing that piracy is a huge moral and financial problem, software developers have adopted and modified another type of system that can be applied to print, audio, video, and streaming media. Called *Digital Rights Management* (DRM), the system was originally devised to protect proprietary information and military information. The idea behind the system is to protect all types of intellectual digital content from anyone who would take it without the consent of the developer(s) or owners. Major companies like Microsoft, Adobe, and IBM are developing and marketing DRM systems, and dozens of smaller companies are springing up.

42.4.1 Purpose. The purpose of DRM is to protect all digital content that originators or owners want protected. DRM permits distributors of electronic content to control viewing access to that content. The content can be text, print, music, or images. Basically, DRM systems use a form of customized encryption. When an end user purchases viewing, listening, or printing rights, an individual "key" is provided. The system works on rules, meaning that although a key is provided, it generally comes with limitations regarding the copying, printing, and redistribution.

Unfortunately, there is no agreement on a DRM solution. Lack of standards is hampering businesses from moving forward with online business initiatives. Because there are so many companies promoting their own incompatible forms of DRM, customers will have to download megabytes of code for each version. Maintaining, upgrading, and managing all of those different versions are major headaches for customers. There

does not appear to be a simple solution; rather than being a technology issue, it is really a matter of business and politics.

42.4.2 Application. Typically, when users become prospective owners of digital rights, they download a content file. The DRM software does an identity check of users, contacts a financial clearinghouse to arrange for the payment to be made, and then decrypts the requested file and assigns users a key. The key is used for future access to the content.

Because the system works on rules, it is possible to impose restrictions. One user might pay just to view material, while another user might want to have printing privileges. A third user might want to download the content to his or her own machine, and, finally, a fourth user might want to have viewing privileges for a specified time. The four different authorized users would thus use the same content, and each would pay according to a rate scale established by the content distributor. Throughout all of the transactions, each user would need a mechanism that allows secure transmissions and identifies that user and the associated level of access privileges.

Although this approach to publishing may sound fairly simple, it is really quite complex. In addition to arranging for different users to access material according to the set rules and to pay according to a rate schedule, it is also necessary for content distributors to handle the back end of the application. Everyone involved in the creation, production, and distribution of the content has to be paid fairly for the use of the content.

Payment is especially important as more and more content providers digitize materials that they can show or print on demand. Many users will read books online, but some physical printing on paper will continue. However, publishers will be able to print precisely those volumes that are requested. This approach will provide customized printing (e.g., large-print editions) as well as saving paper and physical warehouse storage space.

42.4.3 Examples. Several different types of DRM systems exist. Experts agree that the best DRM systems combine both hardware and software access mechanisms. With the advent of the eBook, the digital pad, PDA modems, Internet access devices, and increasingly smaller laptop computers, tying access rights directly to storage media gives publishers and distributors control of where the content is being used as well as by whom.

With the passage of the *Electronic Signatures in Global and National Commerce Act*, referred to as the *E-Sign Bill* and the increasing use of digital signatures, original documents (e.g., legal, medical, or financial) will be stored digitally. President Clinton signed the E-Sign Bill on June 30, 2000, in Philadelphia's Congress Hall using a ceremonial pen and a digital smart card. The bill went into effect on October 1, 2000. The E-Sign Bill gives an online signature the same legal status as a signature etched on paper and makes the digital document the original. Any printout will be considered a copy, so the ability to view documents and videos (e.g., living wills) digitally will actually give the viewer access to the original. Eventually, records for medical treatment and documents for trials may be totally digital and may require submission and viewing digitally. When this becomes the norm rather than the exception, strict adherence to DRM in order to maintain privacy, as well as to provide restitution, will be paramount. In addition, such content-protection schemes will prevent unauthorized modifications of the digital data that would otherwise contribute to fraud.

For example, IBM has released antipiracy technology called the *Electronic Media Management System* that allows for downloading music tracks but puts controls on

how many copies can be made or allows for a copy length limitation to be inserted. Thus, a minute of music could be downloaded to give a listener a taste, but not a chance to pirate the entire music track. To obtain the entire track, the user would be required to pay a fee.

Microsoft distributes free software that embeds metatags in each audio file. The metatags refer back to a central server in which the business rules are stored. This approach requires that material be tagged as it is created; otherwise, if it is released without the embedded tags, it can be illegally copied.

Major companies like Xerox, Microsoft, IBM, and Adobe got heavily involved producing and using this software in the 1990s, and many smaller firms opened shop. As with other new technology launches, eventually many of the small entrants went out of business or were bought by larger firms. Some of the small companies, such as ContentGuard,[21] continue to exist independently, as of this writing.

42.5 PRIVACY-ENHANCING TECHNOLOGIES. Although DRM may seem to be a valid solution for the piracy problem, dissenters feel that DRM and other antipiracy measures give producers and distributors too much control. The rationale is that excessively restrictive rights management may undermine the fair use rights of consumers and academics. Partly as a result of these conflicting viewpoints, many consumers are making increasing use of privacy-enhancing technologies (PET), a broad term for a range of technologies designed to hide the identity and activities of individual users and computers as their traffic traverses the Internet.

42.5.1 Network Proxy. One broad class of tactics and tools used to enhance privacy is based on the concept of a network proxy. In its most general form, a proxy takes a network connection request from a client and redirects it to the ultimate destination, changing the address headers to make it look like the original request came from the proxy itself. When the destination server responds, the proxy returns the results to the requesting client. Proxies have long been used within organizations, both to protect internal users as they make requests to untrusted networks and to track and sometimes block access to undesirable content. For more details on the use of proxies in Web content filtering and monitoring, see Chapter 31 in this *Handbook*.

More recently, proxies have been employed outside the bounds of corporate networks to allow anonymous connections across the Internet. These so-called anonymizing proxies sometimes use encryption to hide the traffic in transit and so also provide protection from traffic analysis. Anonymizing proxies are a serious threat to organizations that desire (or are required by law) to monitor and block users from accessing certain kinds of information. More advanced versions of this concept use multiple routers to hide the path that a request takes through the public network. As a general class, these are known as mixing networks. One example described as early as 1981 is the Chaum Mix.[22] Recently, a concept known as onion routing has become popular in its incarnation as Tor (the onion router), which uses nested layers of traffic encryption as a session travels from one router to the next.[23] See Chapter 70 in this *Handbook* for further information about anonymity on the Internet.

42.5.2 Hidden Operating Systems. Rather than relying on network technologies, some users are choosing to make their actions on the network private by using hidden operating systems. Two basic approaches are common: the virtual machine and the bootable system.

A virtual machine is a system that runs within another operating system. Long used for cross-platform compatibility and the ability to run multiple systems on a single hardware platform, virtual machines, such as Java Virtual Machine, VirtualPC, and VMware, are also used to hide activity from those who would disapprove of it (system administrators, parents, law enforcement, etc.), since the activities of the virtual machine can be made invisible to the host machine.

The bootable system approach, however, stores an entire operating system on some sort of bootable medium, such as a CD or USB device. If a computer can boot from such media and store downloaded content to a peripheral device rather than the host operating system's hard drive, then no record of the usage will remain when the host is next booted up.[24]

42.6 POLITICAL AND TECHNICAL OPPOSITION TO DRM. Vocal opponents of DRM have organized using the Web to exert pressure on vendors using the techniques; criminal hackers and others have developed and circulated software for disabling DRM.

42.6.1 Political Opposition. The *Electronic Frontier Foundation* (EFF) is a widely respected organization, which describes itself as follows:

> From the Internet to the iPod, technologies are transforming our society and empowering us as speakers, citizens, creators, and consumers. When our freedoms in the networked world come under attack, the Electronic Frontier Foundation (EFF) is the first line of defense. EFF broke new ground when it was founded in 1990—well before the Internet was on most people's radar—and continues to confront cutting-edge issues defending free speech, privacy, innovation, and consumer rights today. From the beginning, EFF has championed the public interest in every critical battle affecting digital rights.
>
> Blending the expertise of lawyers, policy analysts, activists, and technologists, EFF achieves significant victories on behalf of consumers and the general public. EFF fights for freedom primarily in the courts, bringing and defending lawsuits even when that means taking on the US government or large corporations. By mobilizing more than 140,000 concerned citizens through our Action Center, EFF beats back bad legislation. In addition to advising policymakers, EFF educates the press and public.[25]

The EFF has steadfastly argued against DRM:

> Digital Rights Management (DRM) technologies attempt to control what you can and can't do with the media and hardware you've purchased.
>
> - Bought an ebook from Amazon but can't read it on your ebook reader of choice? That's DRM.
> - Bought a DVD or Blu-Ray but can't copy the video onto your portable media player? That's DRM.
> - Bought a video-game but can't play it today because the manufacturer's "authentication servers" are off-line? That's DRM.
> - Bought a smart-phone but can't use the applications or the service provider you want on it? That's DRM.
>
> Corporations claim that DRM is necessary to fight copyright infringement online and keep consumers safe from viruses. But there's no evidence that DRM helps fight either of those. Instead DRM helps big business stifle innovation and competition by making it easy to quash "unauthorized" uses of media and technology.

DRM has proliferated thanks to the Digital Millennium Copyright Act of 1998 (DMCA) which sought to outlaw any attempt to bypass DRM.

Fans shouldn't be treated like criminals and companies shouldn't get an automatic veto over user choice and innovation. EFF has led the effort to free the iPhone and other smart phones, is working to uncover and explain the restrictions around new hardware and software, has fought for the right to make copies of DVDs, and sued Sony-BMG for their "rootkit" CD copy protection scheme.[26]

The Free Software Foundation created the *Defective by Design* campaign in May 2006 and organized the first International Day Against DRM. The eighth event took place on May 3, 2013. A recent campaign is *Stop DRM in HTML5*, which is described as follows:

The World Wide Web Consortium (W3C) is considering a proposal to weave Digital Restrictions Management (DRM) into HTML5— in other words, into the very fabric of the Web. Millions of Internet users came together to defeat SOPA/PIPA, but now Big Media moguls are going through non-governmental channels to try to sneak digital restrictions into every interaction we have online. Giants like Netflix, Google, Microsoft, and the BBC are all rallying behind this disastrous proposal, which flies in the face of the W3C's mission to "lead the World Wide Web to its full potential."[27]

Defective by Design is

… [A] participatory and grassroots campaign exposing DRM-encumbered devices and media for what they really are: Defective by Design. We are working together to eliminate DRM as a threat to innovation in media, the privacy of readers, and freedom for computer users. Our actions involve identifying and targeting defective products, pressuring media retailers and hardware manufacturers to stop supporting DRM, exposing the immense concentration of power over media created by DRM, and raising awareness of DRM to libraries, schools, and individuals around the world.[28]

Industry supporters of DRM refer to it as "digital rights management" as if they are the ultimate authority to grant us our rights, as if they are the ones who should have complete and total control over how we use and interact with our media. What they are really doing is managing the restrictions they impose on our media and devices that we would normally have control over in the absence of DRM. We should own our media, not be at the mercy of media companies. For that reason, we refer to it as "Digital Restrictions Management."[29]

Defective by Design argues that DRM is grossly intrusive and unfair:

Amazon's new movie download service is called Unbox and it outlines what DRM implies. The user agreement requires that you allow Unbox DRM software to monitor your hard drive and to report activity to Amazon. These reports would thus include a list of: all the software installed; all the music and video you have; all your computer's interaction with other devices. You will surrender your freedom to such an extent that you will only be able to regain control by removing the software. But if you do remove the software you will also remove all your movies along with it. You are restricted even geographically, and you lose your movies if you ever move out of the USA. You of course have to agree that they can change these terms at any time. Microsoft's newly upgraded Windows Media Player 11 (WMP11) user agreement has a similar set of terms.[29]

Stop DRM Now! argues along the same lines as Defective by Design:

DRM exists for the exclusive benefit of content producers and providers. Companies such as Disney, Sony, and Lion's Gate argue that DRM is needed to prevent people from pirating

music, movies, and other works on P2P networks or by other means. What they won't tell you is that they are really trying to control who, what, when, where, and how you access your music and videos. For instance, when you purchase a song from Apple's iTunes Music Store, you are purchasing a song that can only be played by media applications that support QuickTime or an Apple iPod. Suppose next year you want to buy a new portable music player? Instead of having choice, you must now buy another iPod if you expect to play music you've already purchased.[30]

42.6.2 Technical Countermeasures

42.6.2.1 Reverse Engineering. Reverse engineering allows a programmer to work backward from the finished program or product. Encryption keys can be extracted by reverse engineering playback software. Reverse engineering can circumvent most anti-piracy solutions. As a result, manufacturers of anti-piracy software and hardware are strongly opposed to permitting reverse engineering. The DMCA does allow reverse engineering, but the provisions of DMCA were not intended to enable the circumvention of technical protection measures (TMPs) in order to gain unauthorized access to or to make unauthorized copies of copyrighted works.

42.6.2.2 Published Attacks. The most notable attack on a software key took place when a licensee of CSS neglected to encrypt a decryption key. Obtaining a key by reverse engineering the XingDVD from Xing Technologies, the hackers were then able to guess many other keys. This left the hackers with a collection of decryption keys; even if the XingDVD key was removed, they could still copy DVDs by using the other keys. Using the results of this compromise, a group of people including the Norwegian teenager Jon Lech Johansen developed a program called *DeCSS* to decrypt CSS-encrypted DVDs and play them on Linux machines.

A variety of groups, including the DVD CCA, sued Johansen for publishing tools to subvert copyright protection. Johansen was acquitted twice in Norwegian courts. As of 2007, all remaining lawsuits against him were dropped,[31] and a number of programs like DeCSS are freely available across the Internet. The next section discusses the widespread availability of DRM-cracking tools.

42.6.2.3 Tools for Cracking DRM. Software for breaking DRM controls abounds on the Web. Methods exist to defeat restrictions on:

- CD copying
- DVD copying
- Blu-ray copying
- iTunes copying
- PDF content extraction
- Software activation

We decline to include specific links to such tools in this *Handbook.*

42.7 FUNDAMENTAL PROBLEMS. A number of experts have pointed out that there are fundamental flaws in all the methods for preventing illegal copying of digital materials as described in this chapter. Bruce Schneier, a respected cryptographer, has repeatedly explained that all digital information must be converted to a cleartext

(unencrypted) form before it is displayed or otherwise used. Schneier calls this "the Achilles' heel of all content protection schemes based on encryption."[32] Because the cleartext version has to reside somewhere in volatile or nonvolatile memory for at least some period of time to be usable, it is theoretically possible to obtain a copy of the cleartext version regardless of the complexity of the methods that originally concealed or otherwise limited access to the data. For example, if a DVD movie using complex regional encoding is to be seen on a monitor, at some point the hardware and software that decoded the DVD have to send that data stream to a monitor driver. The decoded data stream is vulnerable to interception. By modifying the low-level routines in the monitor driver, it is possible to divert a copy of the raw data stream to a storage device for unauthorized replay or reconstitution. Similarly, a system may be devised to prevent more than one copy of a document from being printed directly on a printer; however, unless the system prevents screen snapshots, a user can circumvent the restrictions by copying the screen as a bit image and storing that image for later, unauthorized use. Although hardware devices such as dedicated DVD or CD players may successfully interfere with piracy for some time, the problem is exacerbated under the current popular operating systems that have no security kernel and thus allow any processes to access any region of memory without regard to security levels.

42.8 SUMMARY. Piracy is a rapidly growing societal problem affecting a multitude of people and industries. Although producers may suffer the greatest financial losses, there is a substantial impact on consumers. Pirated copies are generally inferior in quality and are sometimes defective. If anything goes wrong, pirated copies, being illegal, are unsupported. Additionally, producers' financial losses due to pirated copies may push the cost of legitimate copies up. Retailers and distributors also suffer due to the loss of sales to pirates. Illegal copies generally are sold more inexpensively than legitimate copies, so retailers and distributors cannot compete on price.

Creative talent, whether software developers, writers, musicians, artists, or performers, plus all the people who helped create the book, magazine, record, performance, painting, concert, or other media, are cheated out of their royalties by pirates. Frequently, because of the amount of time and effort needed to create the end product, the creators depend on the royalties for their livelihood. In addition, poor quality of stolen concepts can irreparably damage the reputation of the creative talent.

Publishers, record companies, art dealers, and other individuals and companies that invest artistic and technical skill along with money and effort to create an original work also lose revenues when that work is pirated. Because of the expenses already laid out to create the original product, companies frequently have to recoup their losses by raising prices for the consumer.

Due to the sophistication of systems and the increased use of the Internet, piracy has become more widespread and has an even greater financial impact worldwide. Many different types of antipiracy systems and techniques have been developed and implemented in an effort to cut down on the ever-increasing instances of piracy. One drawback to all of the systems is the lack of standards applied to software, audio, video, and other media. One of the most promising antipiracy systems is digital rights management. However, even DRM systems are not yet standardized, thus creating even more confusion regarding which is best and what to use.

The media industry is still working out which DRM solution it likes best, or whether it likes DRM at all. In 2007, Apple's Steve Jobs announced that a large portion of EMI's song catalog will be available for DRM-free download from the iTunes Music Store, at a price per song just $0.30 higher than the standard $0.99. Previously downloaded songs,

with DRM, will be upgradeable to a DRM-free version for the difference in the two prices.[33] This move echoes Jobs's recent comments encouraging the music industry to move away from DRM. Apple's actions notwithstanding, some consumers and privacy organizations continue fighting what they see as serious threats to individual privacy in nascent DRM efforts. The use of network proxies and hidden operating systems is adding to the difficulty of discovering inappropriate use of content, much less preventing or prosecuting it. As the balance between content producers and consumers works itself out, it seems likely that many more technologies will come and go, and DRM may be the harbinger of things to come or an unfortunate choice on the way to a bold, new business model for the media industry.

42.9 GLOSSARY. Discussions of piracy and privacy often are riddled with a confusing array of terms and acronyms. Some of the most commonly used terms, organizations, and acronyms are listed in this glossary.

AAC—Advanced Audio Coding. A standardized, lossy compression and encoding scheme for digital audio. Most commonly used format for compressing audio CDs for Apple's iPod and iTunes.

ACATS. Advisory Committee on Advanced Television Service.

Anti-Bootleg Statute (Section 2319A). A U.S. federal statute that criminalizes the unauthorized manufacture, distribution, or trafficking in sound recordings and music videos of "live" musical performances.

ATSC. Advanced Television Systems Committee.

ATV. Advanced Television.

Bootleg recordings. The unauthorized recording of a musical broadcast on radio or television, or at a live concert or performance. These recordings are also known as *underground recordings*.

BSA—Business Software Alliance. A consortium of major software developers, including IBM, Microsoft, Novell, Apple, Dell, and Sun Microsystems, that is attempting to stem lost revenues from pirated computer software. BSA educates computer users on software copyrights and fights software piracy. Individual members, such as Microsoft and Adobe, have their own anti-piracy programs in addition to belonging to the BSA.

CEA. Consumer Electronics Association.

CEMA. Consumer Electronics Manufacturers Association.

CSS—Content Scrambling System. A form of data encryption used to discourage reading media files directly from the disc, without a decryption key. Descrambling the video and audio requires a 5-byte, 40-bit key.

DeCSS—Descrambling Content Scrambling System. A utility developed by Norwegian programmers via reverse engineering and posted on the Web. This utility decrypts CSS and allows individuals to make illegal copies of DVD movies.

DFAST. Dynamic Feedback Arrangement Scrambling Technique.

DMAT. Digital Music Access Technology.

DMCA. The Digital Millennium Copyright Act signed into law October 28, 1998. Designed to implement World Intellectual Property Organization (WIPO) treaties

(signed in Geneva in December 1996); the DMCA strengthens the protection of copyrighted materials in digital formats.

DivX. A brand name of products created by DivX, Inc. (formerly DivXNetworks, Inc.), including the DivX Codec. Known for its ability to compress lengthy video segments into small sizes, it has been the center of controversy because of its use in the replication and distribution of copyrighted DVDs. Many newer DVD players are able to play DivX movies.

DRM—Digital Rights Management. Refers to any of several technologies used by publishers or copyright owners to control access to, and usage of, digital data or hardware and to restrictions associated with a specific instance of a digital work or device.

DVD CCA—DVD Copy Control Association. A not-for-profit corporation that owns and licenses CSS. DVD CCA has filed numerous lawsuits against companies and individuals that make pirated copies of films.

EFF—Electronic Frontier Foundation. A nonprofit organization working in the public interest to protect fundamental civil liberties, including privacy where computers and the Internet are concerned. The organization frequently disagrees with the steps that other organizations and corporations want to take to protect copyrighted materials.

EPIC—Electronic Privacy Information Center. A public interest research group established in 1994 to focus public attention on emerging civil liberties issues and to protect privacy, the First Amendment, and constitutional values.

FairPlay. A digital rights management (DRM) technology created by Apple Inc. (formerly known as Apple Computer), built in to the QuickTime multimedia technology, and used by the iPod, iTunes, and the iTunes Store. Every file bought from the iTunes Store with iTunes is encoded with FairPlay. It digitally encrypts AAC audio files and prevents users from playing these files on unauthorized computers.

FAST—Federation against Software Theft. A group headquartered in Great Britain represents software manufacturers and works with law enforcement agencies in finding and stopping software pirates in Europe.

FCC. Federal Communications Commission.

grpff. A 7-line, 526-character program of Perl code developed by two students at the Massachusetts Institute of Technology. More compact than DeCSS, the program descrambles DVDs but does not contain a decryption key. The code is readily available on the MIT campus via hats, T-shirts, business cards, and bumper stickers.

Hard disk loading. PCs with unlicensed software preinstalled. Use of a single copy of a software program but installed illegally on many machines. The original disks and the documents that should come with the PC are often missing or incomplete.

HDTV—High-definition television. Digital television transmissions are mandated in the U.S. to be the standard.

HRRC—Home Recording Rights Coalition. A coalition representing consumers, retailers, and manufacturers of audio and audiovisual recording products and media. The HRRC dedicates itself to keeping products free of government-imposed charges or restraints on the products' distribution or operation.

IFPI—International Federation of the Phonographic Industry. An organization that promotes the interests of the international recording industry worldwide. Its mission is to fight music piracy; promote fair market access and good copyright laws; help develop the legal conditions and the technologies for the recording industry to prosper in the digital era; and promote the value of music.

IIPA—International Intellectual Property Alliance. A private-sector coalition formed in 1984. The organization represents the U.S. copyright–based industries in efforts to improve international protection of copyrighted materials.

MP3. A technology for downloading music files using the MPEG format via the Internet.

MPAA—Motion Picture Association of America. Composed of member companies that produce and distribute legitimate films and videos. This organization serves as the official voice and advocate of the American motion picture industry. MPAA also assists law enforcement in raids and seizure of pirated videocassettes and DVDs.

MPEG—Moving Pictures Experts Group. A generic means of compactly representing digital video and audio signals for consumer distribution. MPEG video syntax provides an efficient way to represent image sequences in the form of more compact-coded data.

NMPA—National Music Publishers Association. A trade association that represents 700 U.S. businesses that own, protect, and administer copyrights in musical works.

NTSC. National Television System Committee.

PET—Privacy-Enhancing Technologies. The general term for a variety of new technologies and Internet protocols designed to enhance online privacy; includes anonymizing proxies, mixing networks, and onion routing.

Pirated recordings. Unauthorized duplicates of the sounds of one or more legitimate recordings.

RIAA—Recording Industry Association of America. The group has an antipiracy unit that handles initial examination of product on behalf of the recording industry. Pirates can be turned in by calling RIAA at 1-800-BAD-BEAT.

SAG—Screen Actors Guild. The union for screen actors. Members do not get residuals from pirated films, as they do from authorized copies.

SDMI—Secure Digital Music Initiative. A forum of 200 companies from the electronics, music, telecommunications, and information technology industries and the RIAA. The group was active from 1998 to 2001.

SIIA—Software & Information Industry Association. A trade organization of the software and information content industries representing over 800 high-tech companies that develop and market software and electronic content. The organization provides policies and procedures for dealing with software and Internet use within businesses. SIIA also provides guidelines for telling if software is pirated or counterfeited. The SIIA Anti-Piracy Hotline is 1-800-388-7478.

SPA—Software Publishers Association. This association, a division of SIIA, assists in enforcement in dealing with software piracy and also provides education about software piracy. SPAudit Software is one of the first software audit and inventory

tools made available for use by companies (in the 1980s). Improved versions of the software are now available.

Trademark Counterfeiting—Title 18 U.S.C. Section 2320. A federal statute that deals with sound recordings that contain the counterfeit trademark of the legitimate manufacturer or artists.

Trafficking in Counterfeit Labels—Title 18 U.S.C., Section 2318. A federal statute that covers counterfeit labels printed for use on a sound recording.

U.S. Copyright Law (Title 17 U.S.C.). A federal law that protects copyright owners from the unauthorized reproduction or distribution of their work.

WGA—Writers Guild of America. The union for writers of television, video, and film scripts.

WMA—Windows Media Audio. A proprietary compressed audio file format developed by Microsoft Corporation.

42.10 FURTHER READING

Business Software Alliance–USA Homepage. *Anti-Piracy Information*, 2007, www.bsa.org

Cohen, Julie E. "DRM and Privacy." *Berkeley Technology Law Journal*. 2003. www.law.berkeley.edu/institutes/bclt/drm/papers/cohen-drmandprivacy-btlj2003.html

Copyright Act of 1976 (Public Law 94-553); Title 17 U.S.C. Sections 101–120, www4.law.cornell.edu/uscode/html/uscode17/usc_sec_17_00000101-000-.html

Electronic Freedom Foundation home page: www.eff.org

Electronic Privacy Information Center (EPIC) home page: www.epic.org

Gross, T. "The Music Industry, Adapting to a Digital Future: Terry Gross interviews Eliot Van Buskirk." *Fresh Air* podcast, March 13, 2008, www.npr.org/templates/story/story.php?storyId=88145070

Harte, L. *Introduction to Digital Rights Management (DRM); Identifying, Tracking, Authorizing and Restricting Access to Digital Media*. Fuquay Varina, NC: Althos, 2006.

May, C. *Digital Rights Management: The Problem of Expanding Ownership Rights*. Oxford, UK: Chandos Publishing, 2006.

Recording Industry Association of America: www.riaa.com

Schneier.com: www.schneier.com/index.html

Zeng, W., H. Yu, and C-Y. Lin, eds. Multimedia Security Technologies for Digital Rights Management. New York: Academic Press, 2006.

42.11 NOTES

1. *Shadow Market: 2011 BSA Global Software Piracy Study*, 9th ed., May 2012, http://globalstudy.bsa.org/2011/downloads/study_pdf/2011_BSA_Piracy_Study-Standard.pdf
2. www.siia.net
3. "SIIA Shuts Down Notorious Software Pirate, Gaining Large Financial Sum and Cooperation with Further Investigations," April 23, 2012, http://siia.net/index.php?option=com_content&view=article&id=1057:siia-shuts-down-notorious-software-pirate-gaining-large-financial-sum-and-cooperation-with-further-investigations&catid=62:press-room-

4. Microsoft, "Microsoft Incorporates New Anti-Piracy Technologies in Windows 2000, Office 2000," www.microsoft.com/Presspass/press/2000/feb00/apfeaturespr.mspx
5. Microsoft, "Description of the Behavior of Reduced Functionality Mode in Windows Vista," http://support.microsoft.com/kb/925582
6. Microsoft, "Microsoft Product Activation," www.microsoft.com/licensing/resources/vol/default.mspx
7. RIAA, "Anti-Piracy," www.riaa.com/issues/piracy/default.asp (URL inactive).
8. GrayZone Digest, "Worldwide Update," October 1997, www.grayzone.com/1097.htm
9. P. Mercer, "Copycat CDs in an Instant," BBC News, April 16, 2002, http://news.bbc.co.uk/2/hi/entertainment/1930923.stm
10. Recording Industry Association of America, "2005 Commercial Piracy Report," http://76.74.24.142/6BE200AF-5DDA-1C2B-D8BA-4174680FCE66.pdf
11. E. VanBuskirk, "Fans Pay Whatever They Want for Radiohead's Upcoming Album," Underwire Blog from Wired.com, October 1, 2007, http://blog.wired.com/underwire/2007/10/fans-to-determi.html
12. J. Valenti, Napster Statement, February 12, 2001, via Music Industry News Network (mi2n.com), www.mi2n.com/press.php3?press_nb=18419
13. S. Butler, "RIAA Sends Another Wave of Settlement Letters," Billboard.biz, September 20, 2007, www.billboard.biz/bbbiz/content_display/industry/e3i39f76c017d89e0747eaafd53d458f14b (URL inactive).
14. "University of Nebraska Will Bill RIAA $11 for Each Threatening Letter Received," Consumerist.com, March 26, 2007.
15. S. Musil, "Scour to End File-Swapping Service," CNET News.com, November 14, 2000, http://cnet.news.com/2100-1023-248631.html
16. J. Borland, "Broadcasters Win Battle Against iCraveTV.com," CNET News.com, January 28, 2000, http://cnet.news.com/2100-1033-236255.html
17. M. Graziano and Lee Rainie, "The Music Downloading Deluge: 37 Million American adults and youths have retrieved music files on the Internet," April 24, 2001, http://www.pewinternet.org/Reports/2001/The-Music-Downloading-Deluge.aspx
18. J. Silverstein, "iTunes: 1 Billion Served," *ABC News* Website, February 23, 2006, http://abcnews.go.com/Technology/story?id=1653881
19. C. Cohn, "Security Researchers Drop Scientific Censorship Case," www.eff.org/IP/DMCA/Felten_v_RIAA/20020206_eff_felten_pr.html (URL inactive).
20. www.sdmi.org/whats_new.htm
21. www.contentguard.com
22. D. Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Communications of the ACM* 24, No. 2 (February 1981); available online at http://freehaven.net/anonbib/cache/chaum-mix.pdf
23. Tor: Anonymity Online, http://tor.eff.org
24. B. Schneier, "Anonym.OS," www.schneier.com/blog/archives/2006/01/anonymos.html
25. https://www.eff.org/about
26. https://www.eff.org/issues/drm
27. www.fsf.org/campaigns

28. www.defectivebydesign.org
29. www.defectivebydesign.org/what_is_drm_digital_restrictions_management
30. http://stopdrmnow.org
31. C. Cohn, "DVD Descrambling Code Not a Trade Secret," http://www.eff.org/IP/Video/DVDCCA_case/20040122_eff_pr.php (URL inactive).
32. B. Schneier, "The Futility of Digital Copy Prevention," Crypto-Gram Newsletter, May 15, 2001, www.schneier.com/crypto-gram-0105.html
33. American Public Radio, *Marketplace*, April 2, 2007, http://marketplace.publicradio.org/shows/2007/04/02/PM200704024.html (URL inactive).

COMPUTER SECURITY HANDBOOK

COMPUTER SECURITY HANDBOOK

Sixth Edition

Volume 2

Edited by

SEYMOUR BOSWORTH

MICHEL E. KABAY

ERIC WHYNE

WILEY

Cover image: ©iStockphoto.com/Jimmy Anderson
Cover design: Wiley

Copyright © 2014 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Previous Edition: Computer Security Handbook, Fifth Edition. Copyright © 2009 by John Wiley & Sons, Inc. All Rights Reserved. Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permissions.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at http://booksupport.wiley.com. For more information about Wiley products, visit www.wiley.com.

Library of Congress Cataloging-in-Publication Data

Computer security handbook / [edited by] Seymour Bosworth, Michel E. Kabay, Eric Whyne. – Sixth edition.
volumes cm
Includes index.
ISBN 978-1-118-13410-8 (vol. 1 : pbk.) – ISBN 978-1-118-13411-5 (vol. 2 : pbk.) – ISBN 978-1-118-12706-3 (2 volume set : pbk.); ISBN 978-1-118-85174-6 (ebk); ISBN 978-1-118-85179-1 (ebk) 1. Electronic data processing departments–Security measures. I. Bosworth, Seymour. II. Kabay, Michel E. III. Whyne, Eric, 1981–
HF5548.37.C64 2014
658.4′78–dc23

2013041083

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

CONTENTS

PREFACE

ACKNOWLEDGMENTS

A NOTE TO THE INSTRUCTOR

PART I FOUNDATIONS OF COMPUTER SECURITY

1. **Brief History and Mission of Information System Security**
   Seymour Bosworth and Robert V. Jacobson

2. **History of Computer Crime**
   M. E. Kabay

3. **Toward a New Framework for Information Security**
   Donn B. Parker, CISSP

4. **Hardware Elements of Security**
   Sy Bosworth and Stephen Cobb

5. **Data Communications and Information Security**
   Raymond Panko and Eric Fisher

6. **Local Area Network Topologies, Protocols, and Design**
   Gary C. Kessler

7. **Encryption**
   Stephen Cobb and Corinne LeFrançois

8. **Using a Common Language for Computer Security Incident Information**
   John D. Howard

9. **Mathematical Models of Computer Security**
   Matt Bishop

10. **Understanding Studies and Surveys of Computer Crime**
    M. E. Kabay

11. Fundamentals of Intellectual Property Law
William A. Zucker and Scott J. Nathan

PART II THREATS AND VULNERABILITIES

12. The Psychology of Computer Criminals
Q. Campbell and David M. Kennedy

13. The Insider Threat
Gary L. Tagg, CISSP

14. Information Warfare
Seymour Bosworth

15. Penetrating Computer Systems and Networks
Chey Cobb, Stephen Cobb, M. E. Kabay, and Tim Crothers

16. Malicious Code
Robert Guess and Eric Salveggio

17. Mobile Code
Robert Gezelter

18. Denial-of-Service Attacks
Gary C. Kessler

19. Social-Engineering and Low-Tech Attacks
Karthik Raman, Susan Baumes, Kevin Beets, and Carl Ness

20. Spam, Phishing, and Trojans: Attacks Meant to Fool
Stephen Cobb

21. Web-Based Vulnerabilities
Anup K. Ghosh, Kurt Baumgarten, Jennifer Hadley, and Steven Lovaas

22. Physical Threats to the Information Infrastructure
Franklin Platt

PART III PREVENTION: TECHNICAL DEFENSES

23. Protecting the Physical Information Infrastructure
Franklin Platt

24. Operating System Security
William Stallings

25. Local Area Networks
N. Todd Pritsky, Joseph R. Bumblis, and Gary C. Kessler

26. Gateway Security Devices
Justin Opatrny

27. Intrusion Detection and Intrusion Prevention Devices
Rebecca Gurley Bace

28. Identification and Authentication
Ravi Sandhu, Jennifer Hadley, Steven Lovaas, and Nicholas Takacs

29. Biometric Authentication
Eric Salveggio, Steven Lovaas, David R. Lease, and Robert Guess

30. E-Commerce and Web Server Safeguards
Robert Gezelter

31. Web Monitoring and Content Filtering
Steven Lovaas

32. Virtual Private Networks and Secure Remote Access
Justin Opatrny and Carl Ness

33. 802.11 Wireless LAN Security
Gary L. Tagg, CISSP and Jason Sinchak, CISSP

34. Securing VoIP
Christopher Dantos and John Mason

35. Securing P2P, IM, SMS, and Collaboration Tools
Carl Ness

36. Securing Stored Data
David J. Johnson, Nicholas Takacs, Jennifer Hadley, and M. E. Kabay

37. PKI and Certificate Authorities
Santosh Chokhani, Padgett Peterson, and Steven Lovaas

38. Writing Secure Code
Lester E. Nichols, M. E. Kabay, and Timothy Braithwaite

39. Software Development and Quality Assurance
Diane E. Levine, John Mason, and Jennifer Hadley

40. Managing Software Patches and Vulnerabilities
Karen Scarfone, Peter Mell, and Murugiah Souppaya

41. Antivirus Technology
Chey Cobb and Allysa Myers

42. Protecting Digital Rights: Technical Approaches
Robert Guess, Jennifer Hadley, Steven Lovaas, and Diane E. Levine

PART IV PREVENTION: HUMAN FACTORS

43. Ethical Decision Making and High Technology
James Landon Linderman

44. Security Policy Guidelines
M. E. Kabay and Bridgitt Robertson

45. Employment Practices and Policies
M. E. Kabay and Bridgitt Robertson

46. Vulnerability Assessment
Rebecca Gurley Bace and Jason Sinchak

47. Operations Security and Production Controls
M. E. Kabay, Don Holden, and Myles Walsh

48. Email and Internet Use Policies
M. E. Kabay and Nicholas Takacs

49. Implementing a Security-Awareness Program
K. Rudolph

50. Using Social Psychology to Implement Security Policies
M. E. Kabay, Bridgitt Robertson, Mani Akella, and D. T. Lang

51. Security Standards for Products
Paul Brusil and Noel Zakin

PART V DETECTING SECURITY BREACHES

52. Application Controls
Myles Walsh and Susan Baumes

53. Monitoring and Control Systems
Caleb S. Coggins and Diane E. Levine

54. Security Audits
Donald Glass, Richard O. Moore III, Chris Davis, John Mason, David Gursky, James Thomas, Wendy Carr, M. E. Kabay, and Diane Levine

55. Cyber Investigation
Peter Stephenson

PART VI RESPONSE AND REMEDIATION

56. Computer Security Incident Response Teams
Michael Miora, M. E. Kabay, and Bernie Cowens

57. Data Backups and Archives
M. E. Kabay and Don Holden

58. Business Continuity Planning
Michael Miora

59. Disaster Recovery
Michael Miora

60. Insurance Relief
Robert A. Parisi, Jr., John F. Mullen, and Kevin Apollo

61. Working with Law Enforcement
David A. Land

PART VII MANAGEMENT'S ROLE IN SECURITY

62. Quantitative Risk Assessment and Risk Management
Robert V. Jacobson and Susan Baumes

63. Management Responsibilities and Liabilities
Carl Hallberg, M. E. Kabay, Bridgitt Robertson, and Arthur E. Hutt

64. U.S. Legal and Regulatory Security Issues
Timothy Virtue

65. The Role of the CISO
Karen F. Worstell

66. Developing Security Policies
M. E. Kabay and Sean Kelley

67. Developing Classification Policies for Data
Karthik Raman, Kevin Beets, and M. E. Kabay

68. Outsourcing and Security
Kip Boyle, Michael Buglewicz, and Steven Lovaas

PART VIII PUBLIC POLICY AND OTHER CONSIDERATIONS

69. Privacy in Cyberspace: U.S. and European Perspectives
Henry L. Judy, Scott L. David, Benjamin S. Hayes, Jeffrey B. Ritter, Marc Rotenberg, and M. E. Kabay

70. Anonymity and Identity in Cyberspace
M. E. Kabay, Eric Salveggio, Robert Guess, and Russell D. Rosco

71. Healthcare Security and Privacy
Paul Brusil

72. Legal and Policy Issues of Censorship and Content Filtering
Lee Tien, Seth Finkelstein, and Steven Lovaas

73. Expert Witnesses and the *Daubert* Challenge
Chey Cobb

74. Professional Certification and Training in Information Assurance
M. E. Kabay, Christopher Christian, Kevin Henry, and Sondra Schneider

75. The Future of Information Assurance
Jeremy A. Hansen

INDEX

PREFACE

Computers are an integral part of our economic, social, professional, governmental, and military infrastructures. They have become necessities in virtually every area of modern life, but their vulnerability is of increasing concern. Computer-based systems are constantly under threats of inadvertent error and acts of nature, as well as those attributable to unethical, immoral, and criminal activities. It is the purpose of *The Computer Security Handbook* to provide guidance in recognizing these threats, eliminating them where possible and, if not, then reducing any losses attributable to them.

The Handbook will be most valuable to those directly responsible for computer, network, or information security, as well as those who must design, install, and maintain secure systems. It will be equally important to those managers whose operating functions can be affected by breaches in security and to those executives who are responsible for protecting the assets that have been entrusted to them.

With the advent of desktop, laptop, and handheld computers, and with the vast international networks that interconnect them, the nature and extent of threats to computer security have grown almost beyond measure. In order to encompass this unprecedented expansion, *The Computer Security Handbook* has grown apace.

When the first edition of the *Handbook* was published, its entire focus was on mainframe computers, the only type then in widespread use. The second edition recognized the advent of small computers, while the third edition placed increased emphasis on PCs and networks.

Edition	Publication Date	Chapters	Text Pages
First	1973	12	162
Second	1988	19	383
Third	1995	23	571
Fourth	2002	54	1,184
Fifth	2009	77	2,040
Sixth	2014	75	2224

The fourth edition of *The Computer Security Handbook* gave almost equal attention to mainframes and microcomputers, requiring more than twice the number of chapters and pages as the third.

The fifth edition was as great a step forward as the fourth. With 77 chapters and the work of 86 authors, we increased coverage in both breadth and depth. In this sixth edition, we updated all chapters while continuing to cover all 10 domains of the Common Body of Knowledge, as defined by the International Information Systems Security Certification Consortium (ISC)[2]:

1. Security Management Practices: Chapters 10, 12, 13, 14, 15, 19, 10, 31, 43, 44, 45, 46, 47, 48, 49, 50, 51, 54, 55, 62, 63, 64, 65, 66, 67, 68, 74, 75
2. Security Architecture and Models: Chapters 1, 2, 3, 8, 9, 24, 26, 27, 51
3. Access Control Systems and Methodology: Chapters 15, 19, 28, 29, 32
4. Application Development Security: Chapters 13, 19, 21, 30, 38, 39, 52, 53
5. Operations Security: Chapters 13, 14, 15, 19, 21, 24, 36, 40, 47, 53, 57
6. Physical Security: Chapters 4, 13, 15, 19, 22, 23, 28, 29
7. Cryptography: Chapters 7, 32, 37, 42
8. Telecomm, Networks, and Internet Security: Chapters 4, 5, 6, 13, 14, 15, 16, 17, 18, 20, 21, 24, 25, 26, 27, 30, 31, 32, 33, 34, 35, 41, 48
9. Business Continuity Planning: Chapters 22, 23, 56, 57, 58, 59, 60
10. Law, Investigations, and Ethics: Chapters 11, 12, 13, 31, 42, 61

We have continued our practice from the fourth and fifth editions of inviting a security luminary to write the final chapter, "The Future of Information Assurance." We are pleased to include this stellar contribution from Jeremy A. Hansen.

SEYMOUR BOSWORTH
Editor-in-Chief
February 2014

ACKNOWLEDGMENTS

Seymour Bosworth, Editor-in-Chief. I would like to give grateful recognition to Arthur Hutt and Douglas Hoyt, my coeditors of the first, second, and third editions of this *Handbook.* Although both Art and Doug are deceased, their commitment and their competence remain as constant reminders that nothing less than excellence is acceptable. Mich Kabay, my coeditor from the fourth and fifth editions, and Eric Whyne, our fellow editor from the fifth and now sixth editions, continue in that tradition. I would not have wanted to undertake this project without them.

Thanks are also due to our colleagues at John Wiley & Sons: Tim Burgard as former Acquisitions Editor, Helen Cho as Editorial Program Coordinator, Sheck Cho as Executive Editor, Kimberly Kappmeyer as Production Editor, Natasha Andrews as Senior Production Editor, and Darice Moore as Copyeditor. All have performed their duties in an exemplary manner and with unfailing kindness, courtesy, and professionalism.

M. E. Kabay, Technical Editor. I want to thank my beloved wife, Deborah Black, light of my life, for her support and understanding over the years that this project has taken away from our time together. I am also grateful to the authors who have selflessly contributed so much to updating the material presented in this text.

Eric Whyne, Administrative Editor. An undertaking as big as pulling together this handbook would not be possible without my wife Lindsay and the love and support she gives to me and to our son Colton. I'd also like to thank the friends and mentors that have helped me most in my career: Mich and Sy, Tom Aldrich, Tom Payne, Frank Vanecek, and my parents Len and Terri. Any successful undertakings I've had, including this book, have been from listening to the advice they've given and aspiring to internalize the virtues that they exemplify. The authors who have contributed to this book also deserve many thanks for sharing their experience and wisdom. It is something for which I, myself, and the readers are extremely grateful.

A NOTE TO THE INSTRUCTOR

This two-volume text will serve the interests of practitioners and teachers of information assurance. The fourth edition and fifth editions of the *Handbook* were well received in academia; at least one-quarter of all copies were bought by university and college bookstores. The design of this sixth edition continues in the same vein and includes many updates to the material.

University professors looking for texts appropriate for a two-semester sequence of undergraduate courses in information assurance will find the *Handbook* most suitable. In my own work at Norwich University in Vermont, Volume I is the text for our *IS340 Introduction to Information Assurance* and Volume II is the basis for our *IS342 Management of Information Assurance* courses.

The text will also be useful as a resource in graduate courses. In the School of Graduate and Continuing Studies at Norwich University, both volumes have been used as required and supplementary reading for the 18-month, 36-credit Master's of Science in Information Security and Assurance program (MISA).

I will continue to create, update, and post PowerPoint lecture slides based on the chapters of the *Handbook* on my Website for free access by anyone applying them to noncommercial use (e.g., for self-study, for courses in academic institutions, and for unpaid industry training); the materials will be available in the IS340 and IS342 sections:

www.mekabay.com/courses/academic/norwich/is340/index.htm
www.mekabay.com/courses/academic/norwich/is342/index.htm

M. E. Kabay
Technical Editor
October 2013

INTRODUCTION TO PART IV

PREVENTION: HUMAN FACTORS

Human factors underlie all the mechanisms invented by technical experts. Without human awareness, training, education, and motivation, technical defenses inevitably fail. This part details a number of valuable areas of knowledge for security practitioners, including these chapters and topics:

43. **Ethical Decision Making and High Technology.** A strategy for setting a high priority on ethical behavior and a framework for making ethical decisions
44. **Security Policy Guidelines.** Guidelines for how to express security policies effectively
45. **Employment Practices and Policies.** Policy guidelines on hiring, managing, and firing employees
46. **Vulnerability Assessment.** Methods for smoothly integrating vulnerability assessments into the corporate culture
47. **Operations Security and Production Controls.** Running computer operations securely, and controlling production for service levels and quality
48. **E-Mail and Internet Use Policies.** Guidelines for setting expectations about employee use of the Web and e-mail at work
49. **Implementing a Security-Awareness Program.** Methods for ensuring that all employees are aware of security requirements and policies
50. **Using Social Psychology to Implement Security Policies.** Drawing on the science of social psychology for effective implementation of security policies
51. **Security Standards for Products.** Established standards for evaluating the trustworthiness and effectiveness of security products

CHAPTER 43

ETHICAL DECISION MAKING AND HIGH TECHNOLOGY

James Landon Linderman

43.1 INTRODUCTION: THE ABCs OF COMPUTER ETHICS 43·1
- 43.1.1 Why an Ethics Chapter in a Computer Security Handbook? 43·1
- 43.1.2 How Much Time Do You Have for This Chapter? 43·2

43.2 AWARENESS 43·2
- 43.2.1 Principle 1: Ethics Counts 43·2
- 43.2.2 Principle 2: Ethics Is Everybody's Business 43·2
- 43.2.3 A Test: Put Yourself in Another's Shoes 43·2
- 43.2.4 An Approach: Disclose! 43·2

43.3 BASICS 43·3
- 43.3.1 Principle 3: Stakeholders Dictate Ethics 43·3
- 43.3.2 Principle 4: Traditional Principles Still Apply 43·3
- 43.3.3 More Tests 43·3
- 43.3.4 A Guideline Approach: Ask! 43·4
- 43.3.5 Another Guideline Approach: An Ethics Officer 43·4

43.4 CONSIDERATIONS 43·4
- 43.4.1 Principle 5: Ethics Need Not and Should Not Be a Hassle 43·4
- 43.4.2 Principle 6: Ethics Policies Deserve Formality 43·5
- 43.4.3 Principle 7: Ethics Policies Deserve Review 43·5
- 43.4.4 Principle 8: Anticipate 43·6
- 43.4.5 The Smell Test 43·6
- 43.4.6 An Approach: Stocktaking 43·6

43.5 CONCLUDING REMARKS 43·7
- 43.5.1 How to Keep Up 43·7
- 43.5.2 Why to Keep Up 43·7

43.6 FURTHER READING 43·8

43.1 INTRODUCTION: THE ABCs OF COMPUTER ETHICS

43.1.1 Why an Ethics Chapter in a Computer Security Handbook? In an information age, many potential misuses and abuses of information create privacy and security problems. In addition to possible legal issues, ethical issues affect many groups and individuals—including employees and customers, vendors, consultants, bankers, and stockholders—who have enough at stake in the matter to confront and

even destroy an organization over ethical lapses. As is so often the case, consciousness raising is at the heart of maintaining control.

In this chapter, the term "ethics" refers to a system of moral principles that relate to the benefits and harms of particular actions and to the rightness and wrongness of motives and ends of these actions. The major sections cover principles of ethics, tests to help recognize ethical and unethical behavior, and approaches to help ensure good ethical conduct.

43.1.2 How Much Time Do You Have for This Chapter? Section 43.2 requires only one minute to read, but it forms the foundation of this chapter and of a lifelong concern for one of the most important requisites of a valued career and of a civilized society. Ethics matters, and here are the ABCs:

Section 43.2 **A**wareness (a one-minute primer)
Section 43.3 **B**asics (a 10-minute summary)
Section 43.4 **C**onsiderations (a 100-minute study)
Section 43.5 **D**etails (a lifetime of ongoing commitment)

43.2 AWARENESS. The sections that follow distill some of the most important issues.

43.2.1 Principle 1: Ethics Counts. Increasingly, society is holding individuals and organizations to higher ethical standards than in the past; for example, in recent years, many countries have passed privacy legislation, conflict-of-interest restrictions for public officials, and full-disclosure laws for candidates for a variety of public offices. People, individually and collectively, really want to trust others. Two corollaries to the principle that ethics counts are:

1. Good ethical standards are usually good for business.
2. Violations of ethics are almost always bad for business.

In other words, good ethical behavior usually is appreciated by society, and bad ethical behavior almost always is frowned on and punished—sooner or later.

43.2.2 Principle 2: Ethics Is Everybody's Business. A second important principle is that good ethics flourishes best when everyone works at it, both in practicing good ethics and in holding others to do so. The reverse of this is also true: Those who practice bad ethics, or choose to ignore the bad ethics of others, are truly part of the problem. Ethics is inescapably everybody's business.

43.2.3 A Test: Put Yourself in Another's Shoes. One of the best evaluators of whether certain behavior is ethical or not invites you to put yourself in the other person's shoes and ask the role-reversal question: "What if I were on the receiving end of the behavior in question?" This variant of the time-honored golden rule translates to "If I wouldn't like it done to me, I probably shouldn't do it to others."

43.2.4 An Approach: Disclose! One of the best guidelines to help ensure good ethical behavior is to let your stakeholders in on what you are doing or are about to do. Good ethics flourishes in the light of day; bad ethics ultimately relies on concealment.

Disclosure buys you two forms of peace of mind: First, you're being openly honest; second, if others do not like it, you have at least given them the opportunity to express their concerns.

For example, consider organizational policy about managers reading employee email. Almost any policy (anywhere from aggressive intervention to complete hands off) is likely to be ethical if and only if employees are made aware of it.

43.3 BASICS. The expansion of Section 43.2 that follows elaborates on the basic principles enunciated there.

43.3.1 Principle 3: Stakeholders Dictate Ethics. Stakeholders are defined as any individuals or groups with something at stake in the outcome of a decision. In the world of business, stockholders are almost always stakeholders; employees, customers, suppliers, and even competitors are often stakeholders too. How a decision harms or benefits stakeholders is a major ethical consideration. Effects on stakeholders are so important that the best place to start looking at the ethics of a decision is to identify stakeholders and just what it is they have at stake. Decisions will invariably affect individuals and groups, often in opposite ways: Some may stand to gain, others to lose or suffer. The effects and trade-offs raise the principal ethics concerns.

43.3.2 Principle 4: Traditional Principles Still Apply. Recent generations are not the first to raise questions and develop ideas about ethics, although the concept of business ethics has been mocked as an oxymoron and only recently promoted in academia as a valuable area of study. High technology has complicated some issues, but these fundamental principles still apply:

- The Golden Rule ("Do unto others as you would have them do unto you") and its variants have already been mentioned. These principles remain timeless and fundamental.
- Consideration of the interplay of duties, rights, and responsibilities remains important. When making ethical decisions, we normally examine the legal, professional, and customary constraints on behavior that apply to our situation.
- Traditional reasons for good ethical behavior, such as religious principles, egoism, utilitarianism, and altruism, still provide us with a useful taxonomy for discussions about ethics.

The point is that even though modern technology has created new opportunities for unethical behavior and new motivations for good ethical behavior, it has not been necessary to develop new principles to deal with ethics. For example, many of the principles (including politeness) governing the behavior of door-to-door sales still apply to Internet push technology.

43.3.3 More Tests. In Section 43.2.3, we suggested the "other's shoes test" as an excellent evaluator of whether certain behavior is ethical or not. Here we introduce three other tests. The first two are negative in the sense of suggesting that behavior is inappropriate; the third is positive and suggests that the behavior in question is ethical.

1. The "mom test" asks how your mom (spouse, children, best friend, etc.) would react if aware of your actions. If you would be embarrassed having someone close to you know what is going on, the odds are pretty good it is unethical.
2. The "eye-team test" takes this a step further and considers the results of exposing your actions to the whole world as part of an investigative team broadcast. Again, the more embarrassment, the more likely the unethical nature of the actions.
3. The "market test" asks you to think about openly publicizing your actions as a competitive customer relations strategy. Never mind whether such a marketing strategy is actually feasible; if such exposure could impress others favorably, chances are you are on solid ethical ground.

43.3.4 A Guideline Approach: Ask! Section 43.2.4 endorsed disclosure as one of the best guidelines to help ensure good ethical behavior. This simply means letting your stakeholders in on what you are doing, or about to do. Having done so, it then becomes an appropriate guideline to ask those stakeholders for their permission (or acquiescence, or at least acknowledgment) before you proceed. This can be in the form of allowing stakeholders to opt out of certain policies. Many stakeholders prefer an opt-in approach rather than a default assumption of acceptability, particularly if that assumption is nonintuitive or otherwise obscure.

An example of this approach to ethical behavior is the legally enforced requirement in many countries to ask customers for permission to use their personally identifiable information for purposes other than the original defined functions. Typically, customers have either to opt out of information sharing or opt in to such use of their information.

In other cases, stakeholders such as employees or shareholders may disagree with the ethical implications of proposed actions. They can then argue against the proposals to the extent possible. The ultimate option for such stakeholders is to withdraw; employees can resign and shareholders can sell their shares. If the proposed actions are perceived as illegal, critics can become whistle-blowers and report the suspected illegality to law enforcement and regulatory officials.

43.3.5 Another Guideline Approach: An Ethics Officer. Designating an individual to serve as a full- or part-time ethics officer in an enterprise is a powerful, proactive way to help ensure good ethical behavior. Many large organizations now consider this a position on the top management team. But even small organizations can formally delegate such responsibilities on a part-time basis; they need not require much time and energy of the individual involved. In all cases, the common objectives include:

- Clear management commitment to good business ethics, including adequate resources to support this position
- Organizational recognition that this individual has appropriate authority and responsibility, and is a conduit of information into and within the organization
- Hassle-free avenues of access to this person

43.4 CONSIDERATIONS. This section discusses some issues of management style in promulgating and enforcing ethics.

43.4.1 Principle 5: Ethics Need Not and Should Not Be a Hassle. The last thing one wants with business ethics is hassle. An organization's ethics policies

should not be obscure, complicated, onerous, or an obstacle to getting things done. The keys to avoiding hassles include:

- Clear and straightforward articulation of ethics policies
- Consciousness raising as a prime objective of ethics policies
- Clear, comfortable, and safe access to interpretation, for example, by an ethics officer
- Consistency in promulgation, education, application, and enforcement

Employees should not have to guess what constitutes acceptable behavior, nor should they be encumbered by anxiety, guilt, or fear. Individuals should be able to gain clarification on matters before, during, or after events without delay and without fear of being considered a nuisance. Questions about ethics issues deserve unbiased answers, without the presumption that there has been a breach of ethical conduct or that ulterior motives are involved.

Whenever an individual is uncomfortable with a situation or a potential situation, whether he or she is directly involved or not, and particularly if "whistle-blowing" implicates others or even organizational policy, that discomfort needs to be addressed. Even if the individual is wrong and should not be concerned, the discomfort should be dispelled. If the individual is right and there is a legitimate concern, the organization should resolve the issue in a way that does not put the whistle-blower on the spot but rather indicates support for such disclosure.

43.4.2 Principle 6: Ethics Policies Deserve Formality. Like other important policies, an organization's ethics policies deserve formality:

- Clear documentation
- Clear motivation
- Clear sanctions
- Clear management support at every level, including the top

Anything less than the foregoing suggests confusion at best and lip service at worst. Formality should not mean bureaucracy or piles of manuals. Consistent with the fifth principle of avoiding hassles, documentation should be brief and clear, and directed at simplifying matters rather than complicating them. The preparation and presentation of policies should reflect a process of thoughtful, high-priority consideration.

A corollary of this principle is peer participation. Policies that ultimately rely on organizational support are best developed and disseminated with peer involvement.

43.4.3 Principle 7: Ethics Policies Deserve Review. Perhaps the only thing as dangerous as ignorance when it comes to policy is complacency. This is as true of ethics policies as it is of any other organizational policy. In particular, to assume everyone in an organization is on board with policy is naïve at best. Review offers a type of preventive maintenance whereby policies and their promulgation are reconsidered with an eye to improvement.

Any organizational policy requires subscription on the part of the members of the organization. Understanding of and compliance with any policy suffers a dangerous tendency to lapse when a policy simply gathers dust. Even an occasional mention and

discussion of ethical issues or principles during meetings can breathe new life into old policy.

Just as a corollary of Principle 6 was peer participation in the formalization of policy, so is peer involvement a corollary of policy review. Such peer review not only facilitates fresh insights into policy, but the process itself represents a powerful educational opportunity.

43.4.4 Principle 8: Anticipate. Few people relish the prospect of a serious breach of organizational ethics, particularly if matters reach the point of embarrassing publicity or even legal action. It is better to contemplate the worst scenarios in advance rather than to deal with them without preparation and after the fact. Wishful thinking often deters an organization from including appropriate issues in formal policy. This must not be permitted to happen.

It is better to address tough issues head on than to take anything for granted. The two best ways of doing so are to:

1. Have comprehensive policies that cover any foreseeable eventuality.
2. Have a full- or part-time ethics officer in place to stay on top of things.

43.4.5 The Smell Test. Sometimes the other tests discussed (other's shoes, mom, eye-team, market) can result in fuzzy or ambiguous analysis. It may be hard to put yourself in someone else's shoes, especially if different individuals would have widely different reactions to your behavior. And sometimes your family and friends, the general public, or your customers could be neutral or divided in their reactions. The so-called smell test does not require quantitative or even qualitative estimations; it simply relies on your intuition as to whether the behavior in question "smells fishy." In other words, if you catch yourself seeking justifications, or feel a bit uncomfortable even thinking about the implications, the ethics may be as bad or poor as they smell.

43.4.6 An Approach: Stocktaking. Where can an organization start with all this if it has not already done so? Things usually start with a concerned individual (you?) doing some stocktaking and consciousness raising. Questions for you and others in your organization to consider follow. If you like the answers, then your organization is ethically aware. If you do not like the answers, your organization must deal with the issues you have uncovered.

- If you felt that a fellow employee was misusing company resources or harassing someone, is it obvious what you should do? Is there someone in the organization you could comfortably talk with?
- Do you have the sense that top management in your organization is aware of ethics issues? Do you have the sense that they care?
- Do you know if your organization monitors employee email or computer usage? How much personal business, such as email and net surfing, is permissible on company time?
- How important is quality to your company's products and services? Are marketing claims consistent with quality?
- What, if any, information does your company capture about customers, suppliers, or employees without their permission? Without even their knowledge? What, if

any, information does your company provide to entities outside the organization about customers, suppliers, or employees without their permission? Without even their knowledge?

- Does your organization fully comply with license arrangements and payments for software and hardware?
- Are customers, suppliers, and employees always treated with dignity and respect? How would a stakeholder (e.g., a new employee) become aware of organizational ethics policies?

For each of these questions, the issue of how one knows must be raised. Ambiguous, confusing, or unreliable knowledge of what is or is not going on should raise red flags of concern. Discomfort with such knowledge may be symptomatic of underlying problems with ethics.

The stocktaking just suggested should translate into action steps for your organization. Here are a few suggested actions to get started.

- Ask some of your peers to consider the same stock-taking questions, and compare their thoughts with yours.
- Itemize and prioritize any concerns.
- Make an appointment with someone fairly high up in management to discuss those concerns. If you are that someone, make an appointment with two or three peers and with subordinates who represent the staff.
- Examine the feasibility of appointing a full- or part-time ethics officer while considering the downside of not having one.
- Ask your internal or external auditors to consider an audit of your ethics policies the next time they do a financial or operational audit.

43.5 CONCLUDING REMARKS. As with security awareness, ethical awareness needs freshness and repetition.

43.5.1 How to Keep Up. One of the best ways for individuals and organizations to keep up with matters of good ethics is to spread the job around. Do not try to shoulder the effort alone; it will likely overwhelm any one individual, and collective thinking is valuable in these matters. Without abrogating individual responsibilities, charging the right person with the job of ethics officer will certainly help to keep the enterprise on an appropriate ethical course. A growing number of periodicals, both professional and of general interest, include articles involving ethics. Reading and then discussing them with your peers can be invaluable. Additionally, there has been an increase in the number of Websites addressing ethical issues. See Section 43.6.

43.5.2 Why to Keep Up. Keeping on top of organizational ethics is important because it is the right thing to do. Contemporary business practice embraces the idea of an extended value chain where entities such as customers and suppliers, traditionally seen as outside organizational boundaries, are now viewed as partners. Strategic alliances are being formed with these entities, and keeping business partners satisfied and confident is now a strategic necessity. Souring the relationship by a breach of ethics is completely inconsistent with sound business practices. Even if customers and suppliers are not formally viewed as business partners, they are still essential to doing

business and are not to be taken for granted. At the very least, given the range of alternative choices available today, unfair exploitation and other forms of unethical practice expose an organization to outright abandonment.

Finally, in the spirit of total quality management, whereby everyone in the enterprise is seen to contribute to its success or failure, your enterprise can be said to be counting on you. The converse may also be true: Your job and your future career may well depend on the success of your enterprise. The good ethics of your organization reflect favorably on you and your colleagues, but bad ethics will have an opposite effect.

43.6 FURTHER READING. A list follows of some books and Websites that will be helpful to readers seeking further discussion of ethical decision making in general, and ethics in business and high technology.

Barger, R. N. "In Search of a Common Rationale for Computer Ethics," 1994; www.nd.edu/~rbarger/common-rat.html

Computer Ethics Institute. www.computerethicsinstitute.org

Cavazos, E., and G. Morin. *Cyberspace and the Law: Your Rights and Duties in the On-Line World*. Cambridge, MA: MIT Press, 1996.

Computer & Information Ethics Resources on Center for Applied Ethics, University of British Columbia. www.ethicsweb.ca/resources

Vance, David (ed.). "Information System Ethics." http://cyberethics.cbi.msstate.edu

Cyber Citizen Partnership. Information Technology Association of America (ITAA) & Department of Justice (DoJ). www.cybercitizenship.org

Cyberangels. www.cyberangels.org

Cyberspacers (kids' site). www.cyberspacers.com/home.html

EpistemeLinks. Philosophy Resources on the Internet—Computer Ethics. http://tinyurl.com/33kduz (URL inactive).

Ess, C. *Digital Media Ethics*. Polity, 2013.

Ethics and Information Technology (journal). www.springer.com/computer/programming/journal/10676

Floridi, L. "Information Ethics: On the Philosophical Foundations of Computer Ethics." Version 2.0, 1998; www.philosophyofinformation.net/publications/pdf/ieotpfoce2.pdf

Forester, T., and P. Morrison. *Computer Ethics: Cautionary Tales and Ethical Dilemmas in Computing*. Cambridge, MA: MIT Press, 1990.

Gaskin, S., and A. Evans. *GO! Ethics in Cyberspace: Getting Started* 2nd Ed. Prentice-Hall, 2013.

Gotterbarn, D. K. W. Miller, J. Impagliazzo, and A. Z. B. A. Bakar. *Computing Ethics: A Multicultural Approach*. Chapman & Hall, 2013.

Institute for Global Ethics. www.globalethics.org

Johnson, D. O. *Computer Ethics*, 3rd ed. New York: Prentice-Hall, 2000.

Kabay, M. E. "Hacker Tips Published in *Wall Street Journal*." *Network World Security Strategies*, August 28, 2007; www.networkworld.com/newsletters/sec/2007/0827sec1.html

Kabay, M. E. "Ethical Decision-Making: Identifying the Ethical Issue," *Network World Security Strategies*, August 30, 2007; www.networkworld.com/newsletters/sec/2007/0827sec2.html

Kabay, M. E. "Ethical Decision-Making: Using Formal and Informal Guidelines," *Network World Security Strategies*, September 4, 2007; www.networkworld.com/newsletters/sec/2007/0903sec1.html

Kabay, M. E. "Ethical Decision-Making: Principles, Rights and Duties, and Intuitive Cues," *Network World Security Strategies*, September 6, 2007; www.networkworld.com/newsletters/sec/2007/0903sec2.html

Kabay, M. E. "Ethics." Section of Website. www.mekabay.com/ethics/index.htm

Kabay, M. E. "Incident Response: Don't Lie," *Network World Security Strategies*, October 23, 20/07; www.networkworld.com/newsletters/sec/2007/1022sec1.html

Kallman, E. A., and J. P. Grillo. *Ethical Decision Making and Information Technology: An Introduction with Cases*, 2nd ed. New York: McGraw-Hill, 1996.

Kizza, M. *Ethical and Social Issues in the Information Age*, 5th ed. Springer, 2013.

Lessig, L., D. Post, and E. Volokh. "Cyberspace Law for Non-Lawyers." Published via email, 1997; www.ssrn.com/update/lsn/cyberspace/csl_lessons.html (URL inactive).

Online Ethics Center for Engineering and Science. http://onlineethics.org

Orlando, J., and M. E. Kabay. "Social Engineering in Penetration Testing: Cases," *Network World Security Strategies*, October 25, 2007; www.networkworld.com/newsletters/2007/1022sec2.html

Pimple, K. *Emerging Pervasive Information and Communication Technologies (PICT): Ethical Challenges, Opportunities and Safeguards*. Springer, 2013.

Project NEThics at the University of Maryland. www.inform.umd.edu/CompRes/NEThics/ethics (URL inactive).

Schumacher, P., and M. E. Kabay. "Social Engineering in Penetration Testing: Intimidation," *Network World Security Strategies*, November 8, 2007; www.networkworld.com/newsletters/sec/2007/1105sec2.html

Schumacher, P., and M. E. Kabay. "Social Engineering in Penetration Testing: Overload and Fascination," *Network World Security Strategies*, November 13, 2007; www.networkworld.com/newsletters/sec/2007/1112sec1.html

Spinello, R. *Cyberethics: Morality and Law in Cyberspace*, 5th ed. Jones & Bartlett Learning, 2013

Tavani, H. T. *Ethics and Technology: Controversies, Questions, and Strategies for Ethical Computing*. 4th ed. Hoboken, NJ: Wiley, 2012.

Thies, C. *Computer Law and Ethics*. Mercury Learning & Information, 2013.

Thinkquest. "Computer Ethics." http://library.thinkquest.org/26658/?tqskip=1

University of British Columbia Centre for Applied Ethics. www.ethics.ubc.ca

Web Wise Kids. www.webwisekids.com

CHAPTER 44

SECURITY POLICY GUIDELINES

M. E. Kabay and Bridgitt Robertson

44.1 INTRODUCTION 44·1

44.2 TERMINOLOGY 44·2
- 44.2.1 Policy 44·2
- 44.2.2 Controls 44·2
- 44.2.3 Standards 44·2
- 44.2.4 Procedures 44·3

44.3 RESOURCES FOR POLICY WRITERS 44·3
- 44.3.1 ISO/IEC 27002: 2005 44·3
- 44.3.2 COBIT 44·5
- 44.3.3 Informal Security Standards 44·5
- 44.3.4 Commercially Available Policy Guides 44·9

44.4 WRITING THE POLICIES 44·10
- 44.4.1 Orientation: Prescriptive and Proscriptive 44·10
- 44.4.2 Writing Style 44·11
- 44.4.3 Reasons 44·11

44.5 ORGANIZING THE POLICIES 44·11
- 44.5.1 Topical Organization 44·11
- 44.5.2 Organizational 44·12

44.6 PRESENTING THE POLICIES 44·12
- 44.6.1 Printed Text 44·12
- 44.6.2 Electronic One-Dimensional Text 44·13
- 44.6.3 Hypertext 44·13

44.7 MAINTAINING POLICIES 44·14
- 44.7.1 Review Process 44·15
- 44.7.2 Announcing Changes 44·15

44.8 SUMMARY 44·15

44.9 FURTHER READING 44·16

44.10 NOTES 44·16

44.1 INTRODUCTION. This chapter reviews principles, topics, and resources for creating effective security policies. It does not propose specific guidelines except as examples. Many of the chapters in this *Handbook* discuss policy; examples include:

- Chapter 23 provides an extensive overview of physical security policies
- Chapter 25 discusses local area network security issues and policies
- Chapter 38 reviews software development policies
- Chapter 39 surveys quality assurance policies
- Chapter 43 discusses ethics
- Chapter 45 provides guidance on employment policies from a security standpoint

- Chapter 47 includes policies for improving operations security and production
- Chapter 48 reviews specific recommendations for email and Internet usage
- Chapter 49 looks at methods for enhancing security awareness
- Chapter 52 offers policies for secure application design
- Chapters 56 through 59 deal with policies for emergency response, backup, and recovery
- Chapter 61 presents policies for working effectively with law enforcement
- Chapter 66 discusses effective methods for developing security policies in specific organizations

44.2 TERMINOLOGY. One of the preeminent leaders in security policy development, Charles Cresson Wood, has emphasized that when developing policy, it helps to segregate information that has different purposes. Specifically, one should create different documents for policy, standards, and procedures.[1]

44.2.1 Policy. The term "policy" is defined as the rules and regulations set by the organization. Policies are laid down by management in compliance with applicable law, industry regulations, and the decisions of enterprise leaders. Policies are mandatory; they are expressed in definite language and require compliance. Failure to conform to policy can result in disciplinary action, termination of employment, and even legal action. Familiar examples of policy include requirements for background checks when hiring employees, the obligation to follow laws governing the duplication of proprietary software, and restrictions on the use of corporate vehicles for private purposes.

Security policy governs how an organization's information is to be protected against breaches of security; examples include policies on identification and authentication, authorization for specific kinds of access to specific data, responsibilities for data protection, limitations on the use of corporate resources for email and Internet access, and restrictions on installation of programs on corporate systems. Policies are the basis for security awareness, training, and education; they are a necessary underpinning for security audits. Without policies, it is impossible to demonstrate due diligence in the protection of corporate assets.

Policies are focused on the desired results, not on the means for achieving those results. The methods for achieving policies are defined in the next sections on controls, standards, and procedures.

44.2.2 Controls. When developing a framework for implementing security policies, controls are the measures used to protect systems against specific threats. For example, a policy might stipulate that all production systems must be protected against unauthorized modification of data by programmers; a specific control that could be named in the policy might be that test data extracted by programmers from the production databases must be anonymized to protect confidential data.

44.2.3 Standards. A standard in computing can be an accepted specification for hardware, software, or human actions. An example of a technical standard is the Transmission Control Protocol/Internet Protocol (TCP/IP) that governs how systems can be interconnected into the Internet.

Standards can be de facto when they are so widely used that new applications routinely respect their conventions; an example is the Hewlett-Packard interface bus

(HP-IB), which became so popular that it was eventually turned into a *de jure* standard when the Institute of Electrical and Electronics Engineers (IEEE) based its formal IEEE-488 standard on the HP-IB. In contrast, the Centronix parallel interface, although equally popular and universally used, remained proprietary.

In a corporate environment, the term "standard" refers to specific technical choices for implementing particular policies. For example, a corporate policy might stipulate that strong identification and authentication, selected by the technical staff, must be used when gaining access to restricted data; the corresponding standard might specify that a particular brand and model of a microprocessor-equipped smart card should be used in satisfying access control restrictions. Typically, standards are of concern to those who must implement policies; not all standards need be made known to all personnel. Standards also must change in response to a changing technical environment; typically standards change much more rapidly than policies.

44.2.4 Procedures. Procedures prescribe how people are to behave in implementing policies. For example, a policy might stipulate that all confidential communications from employees traveling outside the enterprise must be encrypted; the corresponding standard might define the proprietary virtual private network (VPN) software and hardware needed to implement that policy; and the corresponding procedure would explain in detail each step required to initiate a secure connection using that particular VPN.

44.3 RESOURCES FOR POLICY WRITERS. If one is setting out to create policy *de novo* (i.e., without a preexisting policy document), it is critically important to use an existing policy template. Creating policies without guidance from experienced policy writers is a time-consuming, frustrating job that can consume thousands of hours of time, cause dissension within the enterprise, and leave everyone so disgusted that the policies end up turning into *shelfware*: stored, but never used. There are several well-recognized resources for helping policy writers structure their work, avoid pitfalls, and save enormous amounts of time. In the review that follows, readers will find information about these resources:

- ISO 17799
- COBIT®
- CERT-CC documentation
- NSA Security Guidelines
- U.S. Federal Best Security Practices
- RFC 2196
- IT Baseline Protection Manual
- Commercial policy guides

44.3.1 ISO/IEC 27002:2005. An increasingly popular standard for writing and implementing security policies, especially in Europe, is ISO/IEC 27002:2005, which is the current version of ISO/IEC 17799:2005, in turn based on the old BS7799.

The British Standard 7799 (BS7799) originated in the U.K. Department of Trade and Industry as a code of practice; it was formally renamed the BS7799 in February 1995. BS7799 was not adopted quickly in Great Britain because it was not flexible enough, it used a simplistic security model, and there were more pressing issues, such

as the imminent arrival of the Y2K problem.[2] In addition, BS7799 was a proprietary standard for which users had to pay the equivalent of several hundred dollars before accessing the full documentation.

Version 2 of BS7799 was published in May 1999, and that year also saw the establishment of formal certification and accreditation methods. At that point, the International Organization for Standardization (ISO) began the process of defining BS7799 as an international standard; ISO 17799 was published in 1999. In 2005, the standard was renamed as ISO/IEC 17799:2005 in collaboration with the International Electrochemical Commission (IEC). It was then renamed in 2007.[3]

With the increasing interest in security, ISO/IEC 27002:2005 certification has been established as a goal for many organizations throughout the world. Major consultancies have trained their auditing staff for compliance with ISO/IEC 27002:2005; e-commerce is also a driving force behind the push for certification. One possible motivation is the experience of the ISO 9000 (quality) certification process in the 1980s; certification soon became a competitive edge and then a competitive requirement to maintain and develop market share.

In the context of policy development, ISO/IEC 27002:2005 offers a convenient framework to help policy writers structure their project in accordance with an international standard. The abstract from the ISO follows.

> ISO/IEC 27002:2005 comprises ISO/IEC 17799:2005 and ISO/IEC 17799:2005/Cor.1:2007. Its technical content is identical to that of ISO/IEC 17799:2005. ISO/IEC 17799:2005/Cor.1:2007 changes the reference number of the standard from 17799 to 27002.
>
> ISO/IEC 27002:2005 establishes guidelines and general principles for initiating, implementing, maintaining, and improving information security management in an organization. The objectives outlined provide general guidance on the commonly accepted goals of information security management. ISO/IEC 27002:2005 contains best practices of control objectives and controls in the following areas of information security management:
>
> - security policy;
> - organization of information security;
> - asset management;
> - human resources security;
> - physical and environmental security;
> - communications and operations management;
> - access control;
> - information systems acquisition, development and maintenance;
> - information security incident management;
> - business continuity management;
> - compliance.
>
> The control objectives and controls in ISO/IEC 27002:2005 are intended to be implemented to meet the requirements identified by a risk assessment. ISO/IEC 27002:2005 is intended as a common basis and practical guideline for developing organizational security standards and effective security management practices, and to help build confidence in inter-organizational activities.

The full text of ISO/IEC 27002:2005 is available in electronic format or on paper from the ISO. In addition, a variety of guides are available to help organizations to develop ISO/IEC 27002:2005–compliant policies with minimal rewriting.[4]

44.3.2 COBIT. The *Control Objectives for Information and related Technology* (COBIT) provide a business-oriented set of standards for guiding management in the sound use of information technology.[5] COBIT was developed by volunteers working under the aegis of the IT Governance Institute (ITGI), which was itself founded by the Information Systems Audit and Control Association (ISACA).

COBIT is an information technology (IT) governance framework and supporting tool set that allows managers to bridge the gap between control requirements, technical issues, and business risks. COBIT enables clear policy development and good practice for IT control throughout organizations. COBIT was first published by ITGI in April 1996. ITGI's latest update—COBIT 5, published in 2012—emphasizes regulatory compliance, helps organizations to increase the value attained from IT, highlights links between business and IT goals, and simplifies implementation of the COBIT framework. COBIT 5 is a fine-tuning of the COBIT framework and can be used to enhance work already done based on earlier versions of COBIT. When major activities are planned for IT governance initiatives, or when an overhaul of the enterprise control framework is anticipated, it is recommended to start fresh with COBIT 5. COBIT 5 presents activities in a more streamlined and practical manner so continuous improvement in IT governance is easier than ever to achieve.[6]

COBIT Case Studies provide specific examples of how the framework has been applied.[7] The extensive *Frequently Asked Questions* provide detailed guidance for managers beginning their study of the standard.[8] ISACA also provides an interactive, Web-enabled version of COBIT that allows registered users to "construct and download [their] own, personalized version of COBIT for use on the desktop in MS Word or Access format."[9] Different levels of detail are available for visitors, ISACA members, and purchasers. ISACA also offers training courses[10] as shown:

- COBIT 5 Foundation
- COBIT 5 Implementation
- COBIT 5 Assessor Course

44.3.3 Informal Security Standards. In addition to the formal standards just discussed, several sets of guidelines have garnered a degree of acceptance as the basis for exercising due diligence in the protection of information systems. These informal standards include:

- CERT-CC security improvement modules
- Security guidelines handbook from the U.S. National Security Agency (NSA)
- RFC 2196 from the Internet Engineering Task Force
- IT baseline protection manual from the German Information Security Department

44.3.3.1 CERT-CC® Documentation. The Computer Emergency Response Team Coordination Center (CERT-CC)® at the Software Engineering Institute (SEI) of Carnegie Mellon University (CMU) has compiled a series of security improvement modules (www.cert.org/certcc.html) on these topics:

Vulnerability Remediation

- Current Vulnerability Work
- Vulnerability Notes Database

CSIRT Community

- CSIRT Development
- National CSIRTs

Secure Coding

- Secure Coding Project Page
- Secure Coding Standards

Artifact Analysis

In addition, CERT expert Julia H. Allen published an excellent guide in 2001 that has retained its value.[11] Chapter headings for this 480-page text are listed next.

1. The CERT Guide to System and Network Security Practices
2. Securing Network Servers and User Workstations
3. Securing Public Web Servers
4. Deploying Firewalls
5. Setting Up Intrusion Detection and Response Preparation
6. Detecting Signs of Intrusion
7. Responding to Intrusions
8. Appendix A: Security Implementations
9. Appendix B: Practice-Level Policy Considerations

44.3.3.2 NSA Security Guidelines. The National Security Agency (NSA) of the United States has published a freely available *Security Guidelines Handbook*.[12] The preface describes it in this way:

> This handbook is designed to introduce you to some of the basic security principles and procedures with which all NSA employees must comply. It highlights some of your security responsibilities, and provides guidelines for answering questions you may be asked concerning your association with this Agency. Although you will be busy during the forthcoming weeks learning your job, meeting co-workers, and becoming accustomed to a new work environment, you are urged to become familiar with the security information contained in this handbook.

This set of employee policies is tailored to the needs of the high-security NSA, but it provides useful information that all organizations can adapt to their own requirements. According to the table of contents, these topics are included:

Initial Security Responsibilities

- Anonymity
- Answering Questions about Your Employment
- Answering Questions about Your Agency Training
- Verifying Your Employment
- The Agency and Public News Media

General Responsibilities

- Espionage and Terrorism
- Classification

- Need-to-Know
- For Official Use Only
- Prepublication Review
- Personnel Security Responsibilities
- Association with Foreign Nationals
- Correspondence with Foreign Nationals
- Embassy Visits
- Amateur Radio Activities
- Unofficial Foreign Travel
- Membership in Organizations
- Changes in Marital Status/Cohabitation/Names
- Use and Abuse of Drugs
- Physical Security Policies
- The NSA Badge
- Area Control
- Items Treated as Classified
- Prohibited Items
- Exit Inspection
- Removal of Material from NSA Spaces
- External Protection of Classified Information
- Reporting Loss or Disclosure of Classified Information
- Use of Secure and Nonsecure Telephones

Helpful Information

- Security Resources

44.3.3.3 U.S. Federal Best Security Practices. The United States Federal Chief Information Officers (CIO) Council has created a Best Practices Committee that provides extensive free documentation for policy makers.[13] The committee is defined in this way:

> The Best Practices Committee (BPC) is established by the CIO Council Charter to serve as a focal point for promoting information management/information technology (IM/IT) best practices within the federal government. The BPC brings together a team of IT professionals committed to identifying the most successful of IM/IT practices being implemented in industry, government, and academia, and sharing them with agency CIOs as best practices, to be considered for emulation throughout the Federal government. It is about sharing the successes of others, and *not reinventing the wheel*. It is about constantly learning and applying working models to reduce complexity and achieve results. It is also about cost avoidance and sound stewardship of the taxpayers' dollars.

There is an extensive collection of documents freely available in PDF format for downloading.[14] Some of the topics in the collection of particular interest for this chapter's context include:

- Best Practices
- Enterprise Architecture

- IT Security/Privacy
- GAO Reports
- IT-Related Laws and Regulations

44.3.3.4 RFC 2196 (Site Security Handbook). The Internet Engineering Task Force (IETF) has an extensive list of informational documents called Requests for Comments (RFCs) governing all aspects of the Internet.[15] One document of particular value to any organization trying to improve its security practices is the classic *Site Security Handbook*, RFC 2196, edited by B. Fraser of the Software Engineering Institute at Carnegie Mellon University, the same body that hosts the CERT-CC.[16] The *Handbook* has this structure:

Introduction
- Purpose of this Work
- Audience
- Definitions
- Related Work
- Basic Approach
- Risk Assessment

Security Policies
- What Is a Security Policy and Why Have One?
- What Makes a Good Security Policy?
- Keeping the Policy Flexible

Architecture
- Objectives
- Network and Service Configuration
- Firewalls

Security Services and Procedures
- Authentication
- Confidentiality
- Integrity
- Authorization
- Access
- Auditing
- Securing Backups

Security Incident Handling
- Preparing and Planning for Incident Handling
- Notification and Points of Contact
- Identifying an Incident

- Handling an Incident
- Aftermath of an Incident
- Responsibilities

Ongoing Activities
Tools and Locations
Mailing Lists and Other Resources
References

44.3.3.5 IT-Grundschutz Catalogues. The German government's computer security arm, the *Bundesamt für Sicherheit in der Informationstechnik*, has published a useful set of guidelines since 1997, the *IT-Grundschutzhandbuch.* Originally known in English as the *IT Baseline Protection Manual*, the most recent version as of this writing was published in 2005 and is available free in English as the *IT-Grundschutz Catalogues*.[17] The work is freely available in PDF in German, English, Swedish, and Estonian.

In general, each module presents concepts, threats and vulnerabilities, and countermeasures. This work is easy to understand and provides a sound basis for effective information security protection.

44.3.4 Commercially Available Policy Guides. There are several commercially available policy templates that save time when developing new policies or improving existing policies. Three of the particular values are discussed next.

***44.3.4.1 Charles Cresson Wood's* ISPME.** The most widely used commercially available collection of security standards is the work by Charles Cresson Wood, *Information Security Policies Made Easy* (*ISPME*). The text includes a CD-ROM for easy access to the text so that users can avoid tedious retyping of existing materials.

Wood integrates the perspectives of both management and technical staff when making recommendations. He was one of the original promoters of information security as a way to achieve a competitive advantage and a coauthor of the first computer crime investigation manual. He was one of the first to advocate and document integration of information resource management concepts with information security activities, use of head-count ratio analysis to determine appropriate levels of information security staffing, an information security document life cycle for planning and budgeting purposes, and network management tools to achieve consistent and centralized systems security. He has also developed and successfully marketed two unique software packages that automate information security administration activities. In addition, he evaluated and recommended U.S. Government policies on open versus classified cryptographic research for Frank Press, President Carter's technology advisor.

One of the outstanding features of *ISPME* is that Wood explains every policy and sometimes provides opposing policies for use in different environments. His text is not only a set of templates but an excellent basis for teaching security principles by looking at the practice of security.

44.3.4.2 Tom Peltier's Practitioner's Reference. Tom Peltier is the Year 2001 Hall of Fame Award Recipient from the Information Systems Security Association

(www.issa.org). The citation provides the background that explains why Peltier is so highly regarded in the field of security:

> Tom Peltier is in his fifth decade working with computer technology. During this time he has garnered numerous industry honors for sharing his experiences with follow professionals. Because of his work he was given the 1993 Computer Security Institute's (CSI) Lifetime Achievement Award. In 1999, the Information Systems Security Association (ISSA) bestowed on him its Individual Contribution to the Profession Award, and in 2001 he was inducted into the ISSA Hall of Fame. Tom was also awarded the CSI Lifetime Emeritus Membership Award.

Peltier's policy text is *Information Security Policies and Procedures.* He provides you with the tools you need to develop policies, procedures, and standards. He demonstrates the importance of a clear, concise, and well-written security program. His examination of recommended industry best practices illustrates how they can be customized to fit any organization's needs.

44.3.4.3 SANS Resources. The System Administration and Network Security (SANS) Institute is well known for the excellent security resources it makes available to members and the general public.

44.3.4.3.1 Security Essentials Courses. The SANS Security Essentials Courses (www.sans.org) provide a solid foundation for understanding the issues underlying security policies. Level 524, *Security Policy & Awareness* (www.sans.org), highlights this objective:

> This course is designed to offer an individual a comprehensive approach to understanding security awareness and developing security policy. Business needs change, the business environment changes, and critical systems are continually exposed to new and developing vulnerabilities. Security awareness training is an effective business strategy that reduces the overall risk to an organization, therefore minimizing user-related faults and errors that lead to destructive and costly security incidents. Security awareness and policy development and assessment are a never-ending process.

44.3.4.3.2 Free Resources. The SANS Institute offers many free resources. For details, see www.sans.org/free_resources.php?utm_source=web-sans&utm_medium=ImageReplace&utm_content=Main_resource_button_green&utm_campaign=Home Page&ref=3601. Topics include:

- Reading Room: Over 1,600 computer security white papers in over 70 categories
- Top 20: The Twenty Most Critical Internet Security Vulnerabilities
- Newsletters: Latest Security News

44.4 WRITING THE POLICIES. How should one write security policies? Should they be suggestions? Orders? Positive? Negative? This section affirms that policies should be definite, unambiguous, and directive. In addition, all policies should have (preferably optional) explanations for the reasons behind them.

44.4.1 Orientation: Prescriptive and Proscriptive. Security policies should be written with clear indications that all employees are expected to conform to them. Language should be definite and unambiguous (e.g., "All employees must…" or "No employees shall…"). Some policies require people to do something—these

are *prescriptive* (e.g., "Employees must follow the password procedures defined by the Information Protection Group at all times."). Other policies prohibit certain actions—these are *proscriptive* (e.g., "No employee shall make or order illegal copies of proprietary software under any circumstances.").

44.4.2 Writing Style. Each policy should be short. Simple declarative sentences are best; writers should avoid long compound sentences with multiple clauses. Details of implementation are appropriate for standards and procedures, not for policies. Policies can refer users to the appropriate documents for implementation details; for example, "Passwords shall be changed on a schedule defined in the *Security Procedures* from the Information Protection Group."

For more details on developing policy, see Chapter 45 in this *Handbook*.

44.4.3 Reasons. Few people like to be ordered about with arbitrary rules. Trying to impose what appear to be senseless injunctions can generate a tide of rebellion among employees. It is far better to provide explanations of why policies make sense for the particular enterprise; however, such explanations can make the policies tedious to read for more experienced users. A solution is to provide optional explanations. One approach is to summarize policies in one part of the document and then to provide an extensive expansion of all the policies in a separate section or a separate document. Another approach is to use hypertext, as explained in Section 44.6.3 of this chapter.

44.5 ORGANIZING THE POLICIES. Policies are distinct from the sequence in which they are presented. It is useful to have two distinct presentation sequences for policies: topical and organizational.

44.5.1 Topical Organization. Security involves a multitude of details; how one organizes these details depends on the purpose of the policy document. The most common format puts policies in a sequence that corresponds to some reasonable model of how people perceive security. For example, employees can look at security with a rough correspondence to the physical world. Under this model, one might have a policy document with a table of contents that looks like this:

Principles
Organizational Reporting Structure
Physical Security

- Servers
- Workstations
- Portable computers

Hiring, Management, and Firing Data Protection

- Classifying information
- Data access controls
- Encryption
- Countering industrial espionage

Communications Security

- Perimeter controls
- Web usage and content filtering
- Email usage and privacy
- Telephone and fax usage

Software

- Authorized products only
- Proprietary (purchased) software
- Development standards
- Quality assurance and testing

Operating Systems

- Access controls
- Logging

Technical Support

- Service-level agreements
- Helpdesk functions

44.5.2 Organizational. The complete set of policies may be comprehensive, concise, and well written, but they will still likely be a daunting document, especially for nontechnical staff. To avoid distressing employees with huge tomes of incomprehensible materials, it makes sense to create special-purpose documents aimed at particular groups. For example, one could have guides like these:

- General Guide for Protecting Corporate Information Assets
- Guide for Users of Portable Computers
- A Manager's Guide to Security Policies
- Human Resources and Security
- Network Administration Security Policies
- Programmer's Guide to Security and Quality Assurance
- The Operator's Security Responsibilities
- Security and the Helpdesk

Each of these volumes or files can present just enough information to be useful and interesting to the readers without overwhelming them with detail. Each can make reference to the full policy document.

44.6 PRESENTING THE POLICIES. What options do policy makers have for publishing their policies? This section discusses printing them on paper versus publishing them electronically.

44.6.1 Printed Text. Policies are not inherently interesting. Large volumes full of policies quickly become shelfware. Short paper documents, however, are familiar

to people; they can be carried around, or placed at hand for easy reference anywhere. Reference cards, summary sheets, stickers, and posters are some of the printed media that can be useful in security awareness, training, and education programs. Printed text, like its electronic versions, provides the opportunity for typeface and color to be used in clarifying and emphasizing specific ideas. However, printed copies of policies share a universal disadvantage: They are difficult to update.

Updating dozens, hundreds, or thousands of individual copies of policy documents can be such a headache that organizations simply reprint the entire document rather than struggle with updates. Updates on individual sheets require the cooperation of every user to insert the new sheets and remove the old ones; experience teaches that many people simply defer such a task, sometimes indefinitely, and that others have an apparently limited understanding of the sequential nature of page numbers. Badly updated policy guides may be worse than none at all, especially from a legal standpoint. If an employee violates a new policy that has been promulgated verbally, but available manuals fail to reflect that new policy, it may be difficult to justify dismissal for wrongdoing.

44.6.2 Electronic One-Dimensional Text. Despite the familiarity and ubiquity of paper, in today's world of near-universal access to computers in the work environment, there is a place for electronic documentation of policies. Such publication has enormous advantages from an administrative standpoint. All access to the policies can be controlled centrally, at least in theory. Making the current version of the policies (and subsets of the policies, as explained in Section 44.5.2) available for reference on a server obviates the problem of updating countless independent copies and avoids the normal situation when using paper: chaotic differences among copies of different ages.

How can one cope with employees stubbornly determined to have their own local copies of the policies on their workstations? One solution to this problem of enforcing a single version is to alert every user to changes in the central copy, or to send every user copies of the appropriate documents by email, with a request to replace their copies of lower version number. Although this solution is not perfect, it does help to keep most people up to date.

44.6.3 Hypertext. Perhaps the most valuable contribution from electronic publication of policies is the availability of hypertext. Hypertext allows a reader to jump to a different section of text and then come back to the original place easily. On paper, forward and backward references are cumbersome, and most readers do not follow such links unless they are particularly keen on the extra information promised in the reference. In electronic files, however, additional information may be as easy to obtain as placing the cursor over a link and clicking.

The most important function of hypertext for policy documents is to provide definitions of technical terms and explanations of the reasons for specific policies.

Some users are more comfortable with printed policies. Hypertext, like other formats of text, generally permits users to print out their own copies of all or part of their policy documentation. Many of the tools also allow annotations by users on their own copy of a file.

44.6.3.1 HTML and XML. The most widely used hypertext format today is Hypertext Markup Language (HTML). Its variant Extensible Markup Language (XML), provides additional functionality for programmers, but from the user's perspective, the

hyperlinks are the same. A simple click of the mouse in a Web browser (e.g., Microsoft Internet Explorer, Netscape Communicator, or Firefox) branches to a different page. More sophisticated programming allows the use of frames and, with Java or ActiveX, pop-up windows. Navigation buttons allow the user to move backward to a previous page or forward to another page. Links also can be used to open new windows so that several pages are visible at once. All of these techniques allow users to move freely through a text with full control over the degree of detail they wish to pursue.

44.6.3.2 Rich Text Format and Proprietary Word Processor Files. Some people prefer to use word processor files for hypertext. As long as everyone uses the same word processing software, this approach can work acceptably. For example, it is usually possible to insert a hyperlink to a section of a single document, to a location in a different file on disk, or to a page on the Web. Some word processors, such as Microsoft Word and Corel WordPerfect, allow one to insert pop-up comments; floating the cursor over highlighted text brings up a text box that can provide definitions and commentary.

In addition to explicit links, Microsoft Word and other modern word processing programs can display a table of headings that allows instant movement to any section of the document.

Rich text format (RTF) is a general format for interchanging documents among word processors, but the results are not always comparable. For example, a comment created using Microsoft Word shows up as a pop-up box with a word or phrase highlighted in the text; the same comment and marker read from an RTF file by Corel WordPerfect shows up as a balloon symbol in the left margin of the document.

44.6.3.3 Portable Document Format. Adobe Acrobat's portable document format (PDF) provides all the hyperlinking that HTML offers, but it does so in a form that is universally readable and that can be controlled more easily. The free Acrobat reader is available for multiple operating systems from www.adobe.com. PDF documents can be locked easily, for example, so that no unauthorized changes can be made. In addition, unlike HTML and word processor documents, PDF files can be constructed to provide near-perfect reproduction of their original appearance, even if not all the fonts used by the author are present on the target computer system.

44.6.3.4 Help Files. Help files also provide hypertext capability. In the Windows environment, one can create help files using utilities such as *Help & Manual* from EC Software GmbH. Entering the search string "create help files" into an Internet search engine brings up many pages of such tools. Windows Help files can be distributed easily to any Windows user because they are relatively small, and they are loaded almost instantly by the Help subsystem. In addition, users are permitted to add their own notes to such documents and can easily print out sections if they wish.

44.7 MAINTAINING POLICIES. No fixed policy document can cover all eventualities. The information security field changes constantly and so must policies. Information security is a process much like total quality management: For success, both require a thoroughgoing integration into corporate culture.

Above all, some named individuals must see maintaining security policies as an explicit part of their job descriptions. Hoping that someone will spontaneously maintain security policies is like hoping that someone will spontaneously maintain financial records. However, as explained in Chapter 45 of this *Handbook*, security policies

should represent the best efforts of people from throughout the organization, not the arbitrary dictates of just one person.

44.7.1 Review Process. An information protection working group can meet regularly—quarterly is a good frequency to try—to review all or part of the policies. Employees can be encouraged to suggest improvements in policies or to propose new policies. The working group can identify key areas of greatest change and work on those first, leaving minor policy changes to subcommittees. Members of the working group should discuss ideas with their colleagues from throughout the enterprise, not just with each other. Every effort should contribute to increasing the legitimate sense of involvement in security policy by all employees, including managers and executives.

44.7.2 Announcing Changes. Drafts of the new versions can be circulated to the people principally affected by changes, so that their responses can improve the new edition. Truly respectful inquiry will result in a greater sense of ownership of the policies by employees, although few of them will rejoice in the new policies. Some employees will see new security policies merely as a mild irritant, while others may view them as a tremendous obstacle to productivity and a general nuisance.

Ideally, major changes in policy should be described and explained in several ways. For example, a letter or email (digitally signed, for security) from the president, chair of the board of directors, chief officers, or chief information security officer can announce important changes in policy and the reasons for the changes. A brief article in the organization's internal newsletter, or a spot on the intranet, can also provide channels for communicating the policy decisions to everyone involved.

Finally, the updated policies can be made available or distributed to all employees using some of the channels discussed in Section 44.6.

44.8 SUMMARY. These 10 recommendations will help in preparing to create and implement security policies:

1. Distinguish among policies, controls, standards, and procedures.
2. Use all resources from government, industry bodies, and commercial organizations in preparing to create policies.
3. Use unambiguous prose when defining policies: Tell people what to *do* and what *not* to do.
4. Use short sentences.
5. Give reasons for policies.
6. Provide different views of policies—topical and organizational.
7. Provide several ways of reading the policies, including printed text, electronic text, and hypertext.
8. Review and improve or adapt policies regularly.
9. Circulate drafts showing changes in policies to interested participants before publishing them.
10. Announce major changes using high-level authorities within the enterprise.

44.9 FURTHER READING

Allen, J. H. *The CERT® Guide to System and Network Security Practices*. Reading, MA: Addison-Wesley, 2001.

Atkinson, J. M. "NSA Security Guidelines Handbook," National Security Agency Central Security Service, 2002. www.tscm.com/NSAsecmanual1.html

Barman, S. *Writing Information Security Policies*. Indianapolis: New Riders, 2001.

Boran, S. "IT Security Cookbook." 2009. www.boran.com/security/index.html

Clarke, R. "Best Practice Guidelines: Controls over the Security of Personal Information," 1993, www.anu.edu.au/people/Roger.Clarke/DV/PDSecy.html

Dubin, J. *The Little Black Book of Computer Security*, 2nd ed. 29th Street Press, 2008.

Flynn, N. *The E-Policy Handbook: Designing and Implementing Effective E-Mail, Internet, and Software Policies*, 2nd ed. AMACOM, 2009.

Greene, S. *Security Program and Policies: Principles and Practices*, 2nd ed. Pearson IT Certification, 2013.

Peltier, T. R. *Information Security Policies and Procedures: A Practitioner's Reference*, 2nd ed. Auerbach, 2004.

Portela, I. M., and F. Almeida, eds. *Organizational, Legal, and Technological Dimensions of Information System Administration*. IGI Global, 2013.

Wood, C. C. *Information Security Policies Made Easy, version 12*. Houston: Information Shield, 2012. www.informationshield.com/ispmemain.htm

44.10 NOTES

1. C. C. Wood, *Information Security Policies Made Easy,* version 12 (Houston: Information Shield, 2012), www.informationshield.com/ispmemain.htm
2. *ISO 17799 Information Security Portal,* www.computersecuritynow.com
3. ISO/IEC 27002:2005, "Information Technology—Security Techniques—Code of Practice for Information Security Management," www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=50297
4. Sites that provide free evaluation versions of policy guides include: "Information Security Policy World," www.information-security-policies-and-standards.com, and "Computer Security Policy Directory," www.computer-security-policies.com/index.htm
5. ISACA, "COBIT 5 Product Family," 2013, www.isaca.org/COBIT/Pages/Product-Family.aspx
6. ISACA, "COBIT 5: A Business Framework for the Governance and Management of Enterprise IT," 2013, www.isaca.org/COBIT/Pages/default.aspx
7. ISACA, "COBIT Recognition," 2013, www.isaca.org/COBIT/Pages/Recognition.aspx
8. ISACA, "COBIT FAQs," 2013, www.isaca.org/Knowledge-Center/cobit/Pages/FAQ.aspx
9. ISACA, COBIT Online: www.isaca.org/Template.cfm?Section=COBIT_Online&Template=/ContentManagement/ContentDisplay.cfm&ContentID=15633
10. ISACA, "COBIT 5 Training," 2013, www.isaca.org/Education/COBIT-Education/Pages/COBIT-Training.aspx
11. J. H. Allen, *The CERT® Guide to System and Network Security Practices* (Reading, MA: Addison-Wesley, 2001).

12. NSA, "NSA Security Guidelines Handbook" (Fort Meade, MD: National Security Agency Central Security Service, 2002), www.tscm.com/NSAsecmanual1.html
13. CIO Council BPC: https://cio.gov/about/committees/management-best-practices-committee (URL inactive).
14. CIO Council Documents: https://cio.gov/resources/document-library
15. IETF RFCs: www.ietf.org/rfc.html
16. B. Fraser, "Site Security Handbook," IETF RFC 2196, 1997; www.ietf.org/rfc/rfc2196.txt?number=2196
17. BSI, "IT-Grundschutz," Bundesamt für Sicherheit in der Informationstechnik (German Federal Office for Information Security), 2005; https://www.bsi.bund.de/EN/Topics/ITGrundschutz/itgrundschutz.html

CHAPTER 45

EMPLOYMENT PRACTICES AND POLICIES

M. E. Kabay and Bridgitt Robertson

45.1 INTRODUCTION 45·1

45.2 HIRING 45·1

45.2.1 Checking Candidate's Background 45·2
45.2.2 Employment Agreements 45·3

45.3 MANAGEMENT 45·3

45.3.1 Identify Opportunities for Abuse 45·3
45.3.2 Access Is Neither a Privilege Nor a Right 45·4
45.3.3 The Indispensable Employee 45·4
45.3.4 Career Advancement 45·6
45.3.5 Vacation Time 45·6
45.3.6 Responding to Changes in Behavior 45·7
45.3.7 Separation of Duties 45·8
45.3.8 The Capability and Responsibility Matrices 45·10
45.3.9 No Unauthorized Security Probes 45·11

45.4 TERMINATION OF EMPLOYMENT 45·12

45.4.1 Resignations 45·12
45.4.2 Firings 45·12

45.5 SUMMARY 45·16

45.6 FURTHER READING 45·17

45.7 NOTES 45·17

45.1 INTRODUCTION. Crime is a human issue, not merely a technological one. True, technology can reduce the incidence of computer crimes, but the fundamental problem is that people can be tempted to take advantage of flaws in our information systems. The most spectacular biometric access control in the world will not stop someone from getting into the computer room if the janitor believes it is "just to pick up a listing."

People are the key to effective information security, and disaffected employees and angry ex-employees are important threats according to many current studies. Chapter 13 in this *Handbook* provides detailed information about insider crime.

This chapter presents principles for integrating human resources (HR) management and information security into corporate culture.[1]

45.2 HIRING. The quality of employees is the foundation of success for all enterprises; it is also the basis for effective information security.

45.2.1 Checking Candidate's Background. Hiring new employees poses a particular problem; growing evidence suggests that many people inflate their résumés with unfounded claims. According to Edward Andler, author of *The Complete Reference Checking Handbook*, "cheating on résumés has become distressingly common. And many people are getting by with it, which appears to be making others follow suit." His research shows that up to 10 percent "seriously misrepresent" their background or work histories. A research project run by the Port Authority of New York and New Jersey used an advertisement asking for electricians who were expert at using "Sontag Connectors." They received 170 responses claiming such expertise, even though there was no such device.[2]

Reviewers should be especially careful of vague words such as "monitored" and "initiated." During interviews or background checking, HR staff should find out what the candidate did in specific detail, if possible. All references should be followed up, at least to verify that the candidates really worked where the résumé claims they did.

Unfortunately, there is a civil liberties problem when considering someone's criminal record. Once people have suffered the legally mandated punishment for a crime, whether fines, community service, or imprisonment, discriminating against them in hiring may be a violation of their civil rights. Can one exclude convicted felons from any job openings? From job openings similar to areas in which they abused their former employers' trust? Are employers permitted in law to require that prospective employees approve background checks? Can one legally require polygraph tests? Drug tests? Personality tests?

In some jurisdictions, "negligent hiring" that results in harm to third parties is being punished in civil litigation. Imagine, for example, that a firm were to hire an active criminal hacker as a system administrator without adequate background checking and interviews; if the hacker were then to use his position and corporate resources to break into or sabotage another organization's systems, it is reasonable to suppose that the victim could claim damages from the criminal's employer on the basis of negligent hiring. In addition, "negligent retention" could hold an employer liable when an employee, who may pose a risk to coworkers or the public, is not terminated immediately.

Employers should consult their corporate legal staffs to ensure that they know, and exercise, their rights and obligations in the specific legal context of their work.

Even checking references from previous employers is fraught with uncertainty. Employers may hesitate to give bad references for incompetent or unethical employees for fear of lawsuits if their comments become known, or if the employee fails to get a new job. Today, one cannot rely on getting an answer to the simple question "Would you rehire this employee?"

Ex-employers must also be careful not to inflate their evaluation of an ex-employee. Sterling praise for a scoundrel could lead to a lawsuit from the disgruntled new employer.

For these reasons, a growing number of employers have corporate policies that forbid discussing a former employee's performance in any way, positive or negative. All one gets from a contact in such cases is "Your candidate did work as an Engineer Class 3 from 1991 to 1992. I am forbidden to provide any further information."

It is commonplace in the security field that some people who have successfully committed crimes have been rewarded by a "golden handshake" (a special payment in return for leaving), sometimes even with positive references. The criminals can then move on to victimize a new employer. However, no one knows how often this takes place.

To work around such distortions, interviewers should question candidates closely about details of their education and work experience. The answers can then be checked for internal consistency and compared with the candidate's written submissions. Liars hate details: It is so much harder to remember which lie to repeat to which person than it is to tell the truth.

There are commercial services specializing in background checking (e.g., Achievement Tec). They provide the necessary forms to allow employers to query credit records and other background information. Companies such as Kroll and Securitas Security Services also conduct extensive background checks.

Another way to conduct an employee background check can be done for free or at modest cost via Internet search engines. By entering someone's name into one of these search engines, there is a good possibility that some aspect of the applicant's life will be retrieved. Of particular interest might be a search of messages from a particular blog to see what information is being disseminated.

Experienced employees should interview the candidate and compare notes in meetings to spot inconsistencies. A director of technical support at a large computer service bureau questioned a new employee who claimed to have worked on a particular platform for several years—but did not know how to log on. Had he chatted with any of the programmers on staff before being hired, his deception would have been discovered quickly enough. Ironically, had he told the truth, he might have been hired anyway.

45.2.2 Employment Agreements. Before allowing new employees to start work, they should sign an employment agreement stipulating that they will not disclose confidential information or trade secrets of their previous employers. Another clause must state that they understand that the new employer is explicitly *not* requesting access to information misappropriated from their previous employer, or from any other source. The Uniform Trade Secrets Act, which is enforced in many jurisdictions in the United States, provides penalties that are up to triple the demonstrated financial damages, plus attorney's fees, caused by such data leakage. One high-profile case involved three employees who were found guilty of stealing and trying to sell Coca-Cola secrets to its rival Pepsi.[3]

45.3 MANAGEMENT. Security is the result of corporate culture; therefore, management practices are critically important for successful information protection. External attacks through Internet connections and damage from malicious software are certainly important threats; nonetheless, insider damage due to errors and omissions as well as through dishonesty or a desire for revenge are still major problems for information security.[4] These problems are compounded when there are collaborative threats involving insiders working with those outside the enterprise.

45.3.1 Identify Opportunities for Abuse. Security managers do not have to be paranoid, they just have to act as if they are paranoid. Managers must treat people with scrupulously fair attention to written policies and procedures. Selective or capricious enforcement of procedures may constitute harassment. If some individuals are permitted to be alone in the printer room as salary checks are printed, while other employees of equivalent rank must be accompanied, the latter can justifiably interpret the inconsistency as an implicit indication of distrust. Such treatment may move certain employees to initiate grievances and civil lawsuits, to lay complaints under criminal statutes for discrimination, or even to commit vengeful acts.

45.3.2 Access Is Neither a Privilege Nor a Right. When management removes access rights to the network server room from a system analyst who has no reason to enter that area, the response may be resentment, sulking, and abuse. People sometimes treat access controls as status symbols; why else would a CEO who has no technical training demand that his access code include the tape library and the wiring closet? Managers can overcome these psychological barriers to better security by introducing a different way of looking at vulnerabilities and access. After identifying an opportunity for a particular employee to use the system in unauthorized ways, one should turn the discussion into a question of protecting the person who has unnecessary access against undue suspicion. For example, an employee having more access to secured files than is required is put at risk. If anything ever did go wrong with the secured files, that employee would be a suspect. There is no need to frame the problem in terms of suspicion and distrust.

With these principles in mind, managers should be alert to such dangers as permitting an employee to remain alone in a sensitive area, allowing unsupervised access to unencrypted backups, or having only one programmer who knows anything about the internals of the accounting package.

As for language, it would be better to stop referring to access *privileges.* The very word connotes superiority and status—the last things management should imply. Access is a function and a responsibility, not a privilege or a right; it should be referred to simply as *access functions* or *access authorizations.*

45.3.3 The Indispensable Employee. In many areas of information processing, redundancy is generally viewed as either a bad thing or an unavoidable but regrettable cost paid for specific advantages. For example, in a database, indexing may require identical fields (items, columns) to be placed in separate files (data sets, tables) for links (views, joins) to be established. However, in managing personnel for better security, redundancy is a requirement. Without shared knowledge, an organization is at constant risk of a breach of availability.

Redundancy in this context means having more than one person who can accomplish a given task. Another way of looking at it is that no knowledge should belong to only one person in an organization. Putting the keys to the kingdom in the hands of one employee invites disaster.

Unique resources always put systems at risk; that is why companies such as Tandem, Stratus, and others have so successfully provided redundant and fault-tolerant computer systems for critical task functions, such as stock exchanges and banking networks. These computer systems and networks have twin processors, channels, memory arrays, disk drives, and controllers. Similarly, a fault-tolerant organization will invest in cross-training of all its personnel. Every task should have at least one other person who knows how to do it—even if less well than the primary resource. This principle does not imply that managers have to create clones of all their employees; it is in fact preferable to have several people who can accomplish various parts of any one person's job. Spreading knowledge throughout the organization makes it possible to reduce the damage caused by absence or unavailability of key people.

It is dangerous to allow a single employee to be the only person who knows about a critical function in the enterprise. Operations will suffer if the key person is away, and the enterprise will certainly suffer if this unique resource person decides to behave in unauthorized and harmful ways. Managers should ask themselves if there is anyone in their department whose absence they dread. Are there any critical yet undocumented procedures for which everyone has to ask a particular person?

A client in a data center operations management class volunteered the following story. There was a programming wizard responsible for maintaining a key production program; unfortunately, he had poor communication skills and preferred to solve problems himself rather than to train and involve his colleagues. "It'll be faster for me to do it myself," he used to say. During one of his rare vacations, something went wrong with "his" production program, shutting down the company's operations. The wizard was in the north woods, out of reach of all communications; the disaster lasted until he returned.

Not only does the organization suffer, but also the indispensable persons suffer from the imbalance of knowledge and skill when no one else knows what they know. Some indispensable employees are dedicated to the welfare of their employer and of their colleagues. They may hesitate to take holidays. If their skills are needed from hour to hour, it becomes more difficult for them to participate in committee meetings. These are the people who wear beepers and cannot sit undisturbed even in a two-hour class. If the indispensable employees' skills affect day-to-day operations, they may find it hard to go to offsite training courses, conferences, and conventions. Despite their suitability for promotion, indispensable people may be delayed in their career change because the organization finds it difficult or expensive to train their replacements. In extreme cases, newly promoted managers may find themselves continuing to perform specialized duties that ought to be done by their staff. Sometimes even a VP of Operations is the only person who can make the changes to a production system that should be performed by a programmer three or four levels down.

A particular kind of indispensability occurs when an employee becomes the *de facto* technical support resource for a particular software package or system. Without authorization from their managers, these employees can find themselves in difficulty. They may be fired because their productivity drops too low according to their job descriptions, which do not include providing undocumented technical support to other people. They may burn out and quit because of overwork and criticism. Or they may cause resentment among their colleagues and neighbors by declining to help them, or by complaining about overwork. Alternatively, they may enjoy the situation, and manage to meet all the demands on their time quite successfully, until others in the information technology department begin to feel threatened, and someone either complains to the higher-ups or begins spreading nasty comments about these unauthorized support technicians.

Looking at this situation from a management point of view, there are problems for the recipients of all this free aid. The longer they persist in getting apparently free help from their unofficial benefactor, the longer they can avoid letting upper management know they need help. Then when the bubble bursts and the expert becomes unavailable, managers are confronted with a sudden demand for unplanned resources. In most organizations, unexpected staffing requirements are difficult to satisfy. Managers have a hard time explaining how it is that they were unable to predict the need and to budget for it.

Sometimes persons continue to be indispensable because of fear that their value to their employers resides in their private knowledge. Such employees resent training others. The best way to change their counterproductive attitude is to set a good example; managers should share knowledge with them and with everyone else in their group. Education should be a normal part of the way everyone in the enterprise works. Managers can encourage cross-training by allocating time for it. Cross-training can be a factor in employee evaluations. Current topics from the trade press and academic journals, for example, can be discussed in a journal club, or at informal, scheduled

meetings, where people take turns presenting the findings from recent research in areas of interest.

Reluctance to explain their jobs to someone else may also mask unauthorized or illegal activity. Take, for example, the case of Lloyd Benjamin Lewis, assistant operations officer at Wells Fargo Bank in Beverly Hills, California. He arranged with a confederate outside the bank to cash fraudulent checks for up to $250,000 each on selected legitimate accounts at Lewis's branch. Using a secret code stolen from another branch, Lewis would scrupulously encode a credit for the exact amount of the theft, thus giving the illusion of correcting a transaction error. Lewis stole $21.3 million from his employer between September 1978 and January 1981, when he was caught by accident. For unknown reasons, a computer program flagged one of his fraudulent transactions so that another employee was notified of an irregularity. It did not take long to discover the fraud, and Lewis was convicted of embezzlement. He was sentenced to five years in a federal prison.[5]

Because Lewis was obliged to be physically present to trap the fraudulent checks as they came through the system, he could not afford to have anyone with him watching what he did. Lewis would have been less than enthusiastic about having to train a backup to do his job. If anyone had been cross-trained, the embezzlement would probably not have continued so long, or have become so serious.

45.3.4 Career Advancement. In a topic related to avoiding indispensability, managers can improve the security climate through accepted principles of good human resources management, such as career advancement for all employees. By promoting individuals to new responsibilities, managers can also increase the number of people with expertise in critical functions. As managers carry out their regular employee performance reviews, they should include discussions of each person's career goals. Here, based on a summary by employment expert Lee Kushner, are some practical questions to discuss with employees as part of their interviews.[67]

1. What are your long-term plans?
2. What are your strengths and weaknesses?
3. What skills do you need to develop?
4. Have you acquired a new skill in the past year?
5. What are your most significant career accomplishments, and will you soon achieve another one?
6. Have you been promoted over the past three years?
7. What investments have you made in your own career?

When managers support individuals' interests and aspirations, they foster a climate of respect and appreciation and concurrently support positive feelings about the organization.

45.3.5 Vacation Time. In the example presented in Section 45.3.3, Lloyd Benjamin Lewis took his unauthorized duties (stealing money from his bank) so seriously that during the entire period of his embezzlement, about 850 days, he was never late, never absent, and never took a single vacation day in over two years. Any data center manager should have been quite alarmed at having an employee who had failed to be absent or late a single day in more than two years. The usual rule in companies is that

unused vacation days can be carried over for only a limited time, and then they expire. This is intended to be an incentive to take vacation time; for normal, honest employees, it probably works fine. For dishonest employees who have to be present to control a scam, losing vacation days is intolerable.

Every employee should be required to take scheduled vacations within a definite—and short—time limit. No exceptions should be permitted. Excessive resistance to taking vacations should be investigated to find out why the employee insists on being at work all the time.

Unfortunately, this suspicious attitude toward perfect attendance can cause problems for the devoted, dedicated, and honest employee. An innocent person can get caught up in a web of suspicion precisely because of exceptional commitment. One may be able to avoid difficulties of this kind by:

1. Making the reasons for the policy well known to all employees so no one feels singled out;
2. Relying on the judgment, discretion, and goodwill of the investigating manager to avoid hurt feelings in their most loyal employees; and
3. Switching such an employee's functions temporarily to see if anything breaks.

45.3.6 Responding to Changes in Behavior. Any kind of unusual behavior can pique the curiosity of a manager. Even more important from a security management standpoint, any consistent change in behavior should stimulate interest. Is a normally punctual person suddenly late, day after day? Did an employee start showing up regularly in hand-tailored suits? Why is a usually charming person snarling obscenities at subordinates these days? What accounts for someone's suddenly working overtime every day, in the absence of any known special project? Is a competent person now producing obvious errors in simple reports? How is it that a formerly complacent staffer is now a demanding and bitter complainer?

With so much of the enterprise's financial affairs controlled by information systems, it is not surprising that sudden wealth may be a clue that someone is committing a computer crime. A participant in an information systems security course reported that an accounting clerk at a government agency in Washington, DC, was arrested for massive embezzlement. The tip-off? He arrived at work one day in a Porsche sports car and boasted of the expensive real estate he was buying in a wealthy area of the capital region—all completely beyond any reasonable estimate of his income.

Not all thieves are that stupid. A healthy curiosity is perfectly justified if you see an employee sporting unusually expensive clothes, driving a sleek car after years with a rust bucket, and chatting pleasantly about the latest trip to Acapulco when that person's salary does not appear to explain such expenditures. Unsolicited inquiries into people's private lives, however, will usually win no friends. There is a delicate line to walk, but ignoring the issue does not make it disappear.

The other kind of change—toward the negative—also may indicate trouble. Why is the system manager looking both dejected and threadbare these days? Is he in the throes of a personal debt crisis? In the grip of a blackmailer? Beset with a family medical emergency? A compulsive gambler on a losing streak? On humane grounds alone, one would want to know what is up in order to help; however, a manager concerned with security would be compelled to investigate. In these days of explosive rage and ready access to weapons, ignoring employees with a dark cloud hovering over their heads may be irresponsible and dangerous.

Any radical change in personality should elicit concern. If the normally relaxed head accountant now has beads of sweat on her forehead whenever you discuss the audit trails, perhaps it is time to look into her work more closely. Why does a good family man begin returning from long lunches with whiskey on his breath? A formerly grim manager now waltzes through the office with a perpetual smile on his face. What happened? Or what is happening?

All of these changes should alert managers to the possibility of changes in the lives of their employees. Although these changes do indeed affect the security of an organization, they also concern managers as human beings who can help other human beings. Mood swings, irritability, depression, euphoria—these can be signs of psychological stress. Is an employee becoming alcoholic? A drug addict? Abused at home? Going through financial difficulties? Having trouble with teenagers? Falling in love with a colleague? Of course managers cannot help everyone, and in some cases, help should involve qualified mental health professionals; but at least everyone can express concern and support in a sensitive and gentle way. Such discussions should take place in private, and without alarming the subject or exciting other employees. At any time, a manager should feel free to involve the HR or personnel department. They will either have a psychologist or trained counselor on staff or be able to provide appropriate help in some other way, such as an employee crisis line.

There are sad cases in which employees have shown signs of stress but have been ignored, with disastrous consequences: suicides, murders, theft, and sabotage. Be alert to the indicators and take action quickly.

Australian human resources expert Laura Stack offers this analysis of signs of extreme stress:

> People don't normally all of a sudden flip out; they give off early warning signals. Luckily, managers can observe signs of stress in employee behaviour, beginning with milder signs and culminating in desk rage. Be observant for the following stress stages:
>
> - Physical stage: Headaches, illness, fatigue.
> - Social stage: Negativity, blaming things on others, missed deadlines, working through lunch.
> - Cerebral stage: Clock-watching, errors in assignments, minor accidents, absentmindedness and indecisiveness.
> - Emotional stage: Anger, sadness, crying, yelling, feelings of being overwhelmed, depression.
> - Spiritual stage: Brooding, crying, wanting to make drastic changes in life, not relating well with people, distancing themselves from personal relationships.[7]

The manager's job in probing behavioral changes is difficult; one must walk the thin and possibly invisible line between laissez-faire uninvolvement, risking lifelong regrets or even prosecution for dereliction of duty, and overt interference in the private affairs of the staff, risking embarrassment and possible prosecution for harassment.

Written policies will help; so will a strong and ongoing working relationship with the HR staff. Making it clear to all employees that managers are available for support, but are also expected to investigate unusual behavior, will also help avoid misunderstandings.

45.3.7 Separation of Duties. The same principles that apply to the control of money should apply to control of data. Tellers at a bank, when someone deposits a large check, will always go to a supervisor and have that person look the check over and

initial the transaction. When bank tellers empty the automatic teller machines at night and fill the cash hoppers, there are always two people present. In most organizations, the person who creates a check is not the person who signs it.

In well-run information systems departments, with good operations security, data entry is distinct from validation and verification.[8] For example, a data entry supervisor can check on the accuracy of data entry but should not be allowed to enter a new transaction without having a direct supervisor check the work. There is no excuse for allowing the data entry supervisor to enter a transaction and then, effectively, to authorize it. What if the entry were in error—or fraudulent? Where would the control be?

In quality assurance for program development, the principles of separation of duty are well established. For example, the person who designs or codes a program must not be the only one to test the design or the code.[9] Test systems are separate from production systems; programmers must not have access to confidential and critical data that are controlled by the production staff. Programmers must not enter the computer room if they have no authorized business there; operators must not modify production programs and batch jobs without authorization.[10]

Managers should consider giving up access to functions that have been delegated to two or more subordinates. Maintaining such access could cause more problems than it solves, but in an emergency, access and control could easily be restored. This attitude exemplifies the concept of separation of duties.

In early 1995, the financial world was rocked by the collapse of the Barings PLC investment banking firm. The Singapore office chief, Nicholas Leeson, was accused of having played the futures market with disastrous consequences.[11] The significant point is that he managed to carry out all the orders without independent overview. Had there been effective separation of duties, the collapse would not have occurred.

Another shocking example occurred when a system administrator at UBS PaineWebber, upset about the poor salary bonus he received, deployed malicious code on the company's network. But before quitting his job, he wrote a program that would delete files and wreak havoc on the company's network. By creating a logic bomb, he was able to impact over 1,000 servers and 17,000 individual workstations. Additionally, buying puts against UBS, he would profit from that attack.

A related approach is called dual control. As an example of dual control, consider the perennial problem of having secret passwords not known to managers who sometimes need emergency access to those passwords. This problem does not generally apply to ordinary users' passwords, which normally can be reset by a security administrator without having to know the old password. This temporary password should be changed to a truly secret string by the user, after a single logon. However, to guard against the absence of the only person who has the root password for a system, possibly because the others are on vacation, it is advisable to arrange for dual-access to backup copies of the password. The principle of dual control dictates that such a copy of the root password should be accessible only if two officers of the organization simultaneously sign for it when taking it out of the secure storage:

- One can store a written copy of the root password in a truly opaque envelope, seal it, sign the seal, tape over the seal with nonremovable tape, and then store the envelope in a corporate safe or strongbox
- The password can be encrypted twice using the public keys of two officers; the dual encryption requires the officers to decrypt the ciphertext in the reverse order of encryption (see Chapter 7 in this *Handbook*).

In conclusion, managers should think about the structure of control over information as they design security policies, so that the safeguards afforded by separation of duties or dual control are present throughout all systems.

45.3.8 The Capability and Responsibility Matrices. One of the tools that can help evaluate and improve the resilience of organizations is the capabilities matrix.[12] One starts by listing (or brainstorming, perhaps using Computer-Aided Consensus™)[13] all the critical functions that the organization requires and all the people in the team. The group then has to decide on a way of rating the capabilities of each person; the figure suggests one way to do it, but in no sense is this suggestion meant to constrain users. The group can come to a consensus on which of the team members can do which tasks at which level of competence and then examine the overall pattern of skills. Exhibit 45.1 shows such a matrix with made-up information.

Problems may be highlighted when no one has high-level capabilities or when too few people can effectively carry out those tasks. Another problem is that some people

CAPABILITIES MATRIX FOR UNTELECOM CORPORATION

							Capabilities & Assignments			
Total points ≤	0	1	3	6	9	12	0	1	2	3
EVALUATION	ACK!!!	DANGER!	URGENT	IMPROVE	GOOD	EXCELLENT	NO CAP.	NOVICE	IS/CAN BE BACKUP	CAN BE IN CHARGE
	PERSONNEL CAPABILITIES								TOTAL	NEED
Task	Albert	Betty	Charlie	Dahfla	Edward	Frannie	Golamo	Hur'dath		
application monitoring	3			3			2	2	10	GOOD
application security vulnerabilities	3			3			2	2	10	GOOD
business impact analysis		3		1				2	6	IMPROVE
computer security incident response team			3		3			2	8	IMPROVE
coordination with corporate counsel		3	1		1	2		2	9	GOOD
coordination with human resources group		3	1		1	2		2	9	GOOD
Database performance	1			3			2	3	9	GOOD
Database support	2			3			2	2	9	GOOD
denial-of-service monitoring				1	3			3	7	IMPROVE
denial-of-service response				1	3			2	6	IMPROVE
detect cyber-attacks					3			3	6	IMPROVE
enterprise antivirus									0	ACK!!!
governance decisions		3	1	1		2		2	9	GOOD
identification and authentication			2		2		2	3	9	GOOD
intellectual property protection				2		2		3	7	IMPROVE
intrusion detection systems					3		2	3	8	IMPROVE
intrusion prevention systems			2		3		2	2	9	GOOD
log management systems	1		2	3	3			2	11	GOOD
manage cyber-attacks		2		1	3	2	1	2	11	GOOD
monitor dashboard			2	2	3			2	9	GOOD
network behavior analysis			1	3	2			3	9	GOOD
network discovery			1	2	2			2	7	IMPROVE
penetration testing			1	1	2			3	7	IMPROVE
quality of service measures	1	3		3	1	2		2	12	EXCELLENT
respond to system alarms		2	3	1	3			3	12	EXCELLENT
risk analysis and management				1		3		3	7	IMPROVE
security awareness				1		3		2	6	IMPROVE
security information and event management			3	2				3	8	IMPROVE
service level agreements	2	3	1	3		3		2	14	EXCELLENT
system firewalls			2		3		3	2	10	GOOD
training				1		3		2	6	IMPROVE
unified threat management					3		2	3	8	IMPROVE
vulnerability assessment			2		3		1	2	8	IMPROVE
Web site assessment					2		1	2	5	URGENT
Web site monitoring					3		1	3	7	IMPROVE
wireless intrusion prevention									0	ACK!!!
TOTAL SCORE	13	22	28	42	55	24	23	81		

EXHIBIT 45.1 Capabilities Matrix For Untelecom Corporation

may have so much knowledge compared with others that the organization is in danger were they to be absent; such cases should lead to training efforts to bring others up to an appropriate level of knowledge and skill. Individuals may also be noticeably lacking in skills; again, training may help solve such problems.

A similar matrix should be developed showing the responsibilities of every team member, with indicators of who has primary responsibility and the names of two or more members who can serve as backups (in prioritized order). It is important that every team member fill out their matrix by themselves, because conflicts may be discovered that way.[14]

The responsibility matrix can identify problems in the management structure. For example

- Some tasks may have no one assigned formally or have nobody assigned primary responsibility, leading to breakdowns in response, especially in emergencies.
- There may be tasks where two or more people believe that they are primary decision makers or leaders for a task, leading to conflicts.
- Some people may be identified with far too many assigned tasks; others may have too few.

45.3.9 No Unauthorized Security Probes. In general, all managers—not just security officers—should always be looking for vulnerabilities and opportunities for improving security. However, no one should ever *test* production systems for vulnerabilities without the full cooperation of the corporate information protection group, and only with authorization of the right executives. Written approval for explicit tests of security are informally known as *get-out-of-jail cards*, because without them, employees can go to jail for unauthorized probes of system security.

The case of Randal Schwartz, a consultant to Intel Corporation in Beaverton, Oregon, is a salutary example for employees of the dangers of unauthorized security probes. He was convicted of hacking his way into Intel Corporation computer networks in what he claimed was an effort to point out security flaws while he was working there as a consultant. The would-be security expert failed to notify his employers of his intentions and forgot to get authorization for stealing passwords and making unauthorized changes in system software. He was convicted of three felony counts in July 1995 and was fined $68,000 in restitution as well as being put under five years of probation and having to perform 480 hours of community service.[15]

A counterexample to warn managers of misplaced zeal in suppressing cooperation with law enforcement is the case of Shawn Carpenter, a network intrusion detection security analyst at Sandia National Laboratories. He was fired by publicity-shy administrators when he worked with law enforcement officials to track down extensive penetrations of U.S. national security assets. An investigation code-named TITAN RAIN began in late 2003.[16] Carpenter noted a flood of expert hacker activity focusing on data theft from a wide range of national security interests. Carpenter discovered that "the attacks emanated from just three Chinese routers that acted as the first connection point from a local network to the Internet."[17] Carpenter worked with U.S. Army Counterintelligence and FBI investigators to learn more about the attacks and the attackers. Carpenter never used Sandia's or government-owned equipment or network resources in his investigations. Administrators applied Sandia Internal Directive 12 ISNL ID012, which "specifically prohibits employees from speaking with local,

state, or Federally elected officials." In 2007, Carpenter was awarded $4.3 million for wrongful termination.[18]

In February 2012, police and courthouse managers in Vienna, Austria, staged a theatrically convincing terrorist attack as a training exercise. The staging included "one simulated death, apparently by a gunshot to the head. Makeup was used to simulate injuries, and several officers were placed in the building as if they were injured persons. The supposed death was staged in front of courthouse staff who were evacuating offices." Because none of the ordinary staff members was informed (and nothing was announced to the public), "By the next day 40 staff members were in treatment for severe trauma and an undisclosed number had taken sick leave. We must assume that some will suffer from post-traumatic stress disorder (PTSD) in the weeks and months to come."[19] See Chapter 46 in this *Handbook* for more discussion of how to prepare for tests of this sort.

45.4 TERMINATION OF EMPLOYMENT. Taking our security mandate in the widest sense, we have to protect our employer and ourselves against potential damage from unethical, disgruntled, or incompetent employees, and against the legal consequences of improper firing procedures. Common sense and common decency argue for humane and sensitive treatment of people being fired and those who are resigning. Firing people is a stressful time for everyone concerned, and it usually leads to increased security risks.[20] Managers should do everything in their power to ensure a courteous, respectful, and supportive experience when terminating employment.

45.4.1 Resignations. Potentially the most dangerous form of employment termination is a resignation. The problem is summed up in the caption of a cartoon where a savage attack is in progress against a medieval town that is in flames; a clan war chieftain confronts a singed and dirty warrior. "No, no, Thor! Pillage, THEN burn!" Like the war chieftain, employees rarely resign without planning. An employee may have an indefinite period during which the action is imminent, while the employer may remain unaware of the situation. If the employee has bad feelings toward, or evil designs on, the current employer, there is a period of vulnerability frequently unknown to management. Dishonest or unbalanced employees could steal information or equipment, cause immediate or delayed damage using programmatic techniques, or introduce faulty data into the system.

The policies discussed in previous sections of this chapter should reduce the risks associated with resignations. The manager's goal should be to make resignations rare and reasonable. By staying in touch with employees' feelings, moods, and morale, managers can identify sources of strain and perhaps resolve problems before they lead to resignations and their associated security risks.

45.4.2 Firings. Firings appear to give the advantage to employers, but there may be complications.

45.4.2.1 Timing. One advantage is that the time of notification to a fired employee can be controlled to minimize effects on the organization and its business. For example, employers might find it best to fire an incompetent, or no longer acceptable, employee before beginning an important new project or after a particular project has finished.

Some people argue that to reduce the psychological impact on other employees, they should fire people at the end of the day, perhaps even before a long weekend.

The theory is that the practice gives everyone a cooling-off period outside working hours. These managers say they do not want the buzz of conversation and speculation that often follow a firing to intrude on the workday. This policy fails to regard the psychological stress to employees who have a ruined weekend and no way of responding constructively to their potentially catastrophic loss of a regular income.

A better approach to this stressful task is to fire people early on Monday morning in order to provide an unrushed exit interview and, if appropriate, job counseling to help the employee prepare for job hunting. In this scenario, the regrettable necessity (from the manager's point of view) of terminating employment is buffered by professionals in the HR department, who can give the departing employee a sense of hope and some practical as well as emotional support in this difficult time. A humane attitude is particularly important during downsizing, or when plants are closed, and many people are being fired—one of the worst experiences possible for both employees and managers, and an event that has serious security implications.

In one large company, the personnel department asked their information security staff to suspend the access codes for more than 100 people who were to be fired at 6:00 PM on Tuesday. On Wednesday at 8:00 AM, the security staff began receiving phone calls asking why the callers' logon IDs no longer worked. It turned out that the personnel staff had failed to inform the terminated employees on time. The psychological trauma to both the employees who were fired and to the security staff was severe. Several security staff members were sent home in tears to recuperate from their unfortunate experience. The harm done to the fired employees was even more serious, and the effect on morale of the remaining employees was a disaster. There could well have been violence in that situation.

45.4.2.2 Procedures upon Termination. In both resignations and firings, security consultants unanimously advise instant action. Not for them the leisurely grace period during which employees wind down their projects, or hand them off to other staff members. Security officers are a hard lot, and they usually advise this scenario: In a formal exit interview, and in the presence of at least two managers, an officer of the employer informs the employee politely that his or her employment is at an end. During the exit interview, the officer explains the reasons for termination of employment. The officer gives the employee a check for the period of notification required by law or by contract, plus any severance pay due. Under supervision, preferably in the presence of at least one security guard, the employee is escorted to the accustomed work area and invited to remove all personal belongings and place them in a container provided by the employer. The employee returns all company badges, IDs, business cards, credit cards, and keys, and is then ushered politely outside the building.

At the same time all this is happening, all security arrangements must be changed to exclude the ex-employee from access to the building and to all information systems. Such restrictions can include:

- Striking the person's name from all security post lists of authorized access
- Explicitly informing guards that the ex-employee may not be allowed into the building, whether unaccompanied or accompanied by an employee, without special authorization by named authorities
- Changing the combinations, reprogramming access card systems, and replacing physical keys if necessary for all secure areas to which the individual used to have authorized access

- Removing or changing all personal access codes known to have been used by the ex-employee on all secured computer systems, including microcomputers, networks, and mainframes
- Informing all outside agencies (e.g., tape storage facilities and outsourced functions) that the ex-employee is no longer authorized to access any of the employer's information or to initiate security or disaster recovery procedures
- Requesting cooperation from outside agencies in informing the employer if ex-employees attempt to exercise unauthorized functions on behalf of their former employer

The task is made more difficult by seniority, or if the ex-employee played an important role in disaster recovery or security. The employer should be assiduous in searching out all possible avenues of entry resulting from the person's position of responsibility and familiarity with security procedures.

In one story circulating in the security literature, an employee was fired without the safeguards just suggested. He returned to the workplace the next Saturday with his station wagon and greeted the security guard with the usual friendliness and confidence. The guard, who had known him for years, was unaware that the man had been fired. The ex-employee still had access codes and copies of keys to secure areas. He entered the unattended computer room, destroyed all the files on the system, and then opened the tape vault. He engaged the guard's help in loading all the company's backup tapes into his station wagon. The thief even complained about how he had to work on weekends. This criminal then tried to extort money from the company by threatening to destroy the backup tapes, but he was found by police and arrested in time to prevent a disaster for his ex-employer.

This story emphasizes the importance of reaching *everyone* who needs to know that an employee no longer works for the enterprise.

45.4.2.3 Support in Involuntary Terminations. Security does sometimes prevent a farewell party, one obvious sign of friendliness. The problem with a farewell party at work is that employees leaving under a cloud may feel humiliated when other people get a party but they do not. Generally, it makes sense to treat all departing employees the same, even if the termination is involuntary.

However, nothing stops a humane and sensitive employer from encouraging employees to arrange an after-hours party even for people who have been fired. If a *resignation* is on good terms, however, the employer may even arrange a celebration, possibly during working hours and perhaps at company cost, without having to worry about possible negative repercussions.

A firing, or a resignation on poor terms, has two psychological dangers: effects on the individual concerned of embarrassment, shame, and anger, and effects on the remaining staff of rumors, resentment, and fear. Both kinds of problems can be minimized by publishing termination procedures in organization documents provided to all employees; by requiring all employees to sign a statement confirming that they have read and agreed to the termination procedures; and by consistent application of the termination procedures.

The personal shock of being fired can be reduced by politeness and consideration consistent with the nature of the reasons for being fired, although even nasty people should not be subject to verbal or physical abuse, no matter how bad their behavior. Their treatment should be consistent with that meted out to other fired employees, and there should be generous severance arrangements, if possible.

Organizational turmoil can be reduced by convening organization-wide or departmental meetings to brief remaining employees on the details of a significant termination. Open discussions, including how people feel about the rupture of relationships, can be helpful. The remaining employees may have to suffer grief, as a process, not a state. Grief is a normal and healthy response to disruption of relationships (e.g., death of a loved one, divorce, and even the loss of a coworker). Some people value social relationships more than other aspects of their work, and they may be especially affected by firings. Grief involves stages of denial, anger, mourning, and recovery. Trying to forestall such responses by denying that people legitimately have feelings is foolish and counterproductive. It is far better to encourage those who are upset to voice their feelings, and to engage in constructive discussion, than to clamp down in a futile attempt to suppress discussion.

45.4.2.4 Style of Termination. The way an organization handles job termination affects more than internal relations; it also influences its image in the outside world. Prospective employees will think twice about accepting job offers from an organization that mistreats departing employees. Clients may form a negative impression of a company's stability if it abuses its own people. Investors also may look askance at a firm that gets a reputation for shoddy treatment of employees. Bad relations among managers and employees are a warning sign of long-term difficulties.

45.4.2.5 Legal Issues. There is another dimension to employment termination that depends on local laws and the litigation environment. The United States, for example, is said to be one of the most litigious nations on the planet, perhaps because of the high number of lawyers compared with the total population.

The list that follows is not legal advice; for legal advice, consult an attorney. However, simple experience does teach some principles, even without going to law school. Here are some pragmatic guidelines for preventing legal problems related to firings for cause:

- Build a solid, documented case for firing someone before acting.
- Keep good records, be objective, and get the opinions of several trustworthy people on record.
- Offer the delinquent employee all reasonable chances to correct his or her behavior.
- Give the employee clear feedback long before considering firing.

Timing is important in employee relations, as it is in almost everything else we do. In particular, if an employee is found to be behaving improperly or illegally, there must be no marked delay in dealing with the problem. Such persons could sue the employer and individual managers. They could argue in court that the very fact that there was a delay in firing them was proof that the firing was due to other factors such as personality conflicts, racism, or sexism. A well-defined procedure for progressing through the decision will minimize such problems.

The critical legal issue is consistency. If rules such as those just described for the day of the firing are applied haphazardly, there could easily be grounds for complaining of unfairness. Those to whom the rules were strictly applied would justifiably feel implicitly criticized. How would we feel if we were singled out by having guards check what we took home from our desk—if everyone else got a party and two weeks' notice? Such inconsistency would be grounds for legal proceedings for defamation of character. The company might lose and it might win, but what nonlawyer wants to spend time in court?

Another issue that arises in connection with firings and resignations involves nondisclosure agreements. All such agreements must be included in a contract signed before the prospective employee begins work; it is almost impossible to force an existing employee to sign such an agreement.

Managers, the legal department, and the personnel department should study the necessity and feasibility of instituting a legally binding contractual obligation to protect their company's confidential information for a specified period of time after leaving. One typically does not impose indefinite gags on people, as one year seems to be normal. (However, there are exceptions. Oprah Winfrey insists that all employees who work at Harpo sign a lifelong confidentiality agreement, which an Illinois appeals court upheld when a former employee tried to write a book about the media mogul.) For this measure to be meaningful, the initial employment contract should stipulate that departing employees must reveal their new employer, if there is one at that time.

Noncompetition agreements require the employee to refrain from working for direct competitors for perhaps a year after termination of employment. The key to a successful clause here is that there be a strict, operational definition of "direct competitors." Because this limitation can be an onerous impediment to earning a living, many jurisdictions forbid such clauses.

45.5 SUMMARY.

Some of the key recommendations from this chapter follow:

Hiring

- Investigate the accuracy of every likely job candidate's résumé.
- Perform background investigations when hiring for sensitive positions.
- Arrange for experienced staff members to interview candidates and discuss inconsistencies.
- Require signing of a legally appropriate employment contract.

Ongoing Management

- Identify and resolve opportunities for abuse.
- Assign access functions on the basis of need, not social status.
- Identify indispensable employees, and arrange for cross-training of other staff.
- Require employees to take their vacations, or to rotate their job functions periodically, so as to assure operational continuity and as a possible indication of fraud.
- Note and respond to sudden changes in behavior and mood; involve human resources as appropriate.
- Enforce separation of duties and dual control for sensitive functions.
- Do not engage in, or tolerate, unauthorized probes of system security.

Termination of Employment

- Provide an opportunity for fired employees to receive counseling and support.
- Ensure that the HR department collaborates with the information technology group to take all appropriate security measures when anyone leaves the employment of the enterprise.

- Ensure that firings do not cause long-term morale problems.
- Follow the guidance of corporate counsel to avoid wrongful dismissal suits.
- Use legally appropriate nondisclosure and noncompetition clauses in employment contracts.

In summary, information security depends on coordination with HR personnel to ensure consistent policies for hiring, ongoing management, and termination of employment.

45.6 FURTHER READING

Armstrong, M. *Armstrong's Handbook of Human Resource Management Practice*, 12th ed. Kogan Page, 2012.

Cascarino, R. E. *Corporate Fraud and Internal Control + Software Demo: A Framework for Prevention*. Wiley, 2012.

Dresang, D. L. *Personnel Management in Government Agencies and Nonprofit Organizations*, 5th ed. Pearson, 2009.

Girgenti, R. H., and T. P. Hedley. *Managing the Risk of Fraud and Misconduct: Meeting the Challenges of a Global, Regulated and Digital Environment*. McGraw-Hill, 2011.

Mathis, R. L., and J. H. Jackson. *Human Resources Management*, 13th ed. South-Western Cengage Learning, 2010.

McNamara, C. "All About Human Resources and TalentManagement." 2013, http://managementhelp.org/humanresources/index.htm

NAPA (National Academy of Public Administration). "Browse Publications by Category." 2013. www.napawash.org/publications-reports/browse-publications-by-keyword

NOLO. "Human Resources." 2013. www.nolo.com/info/human-resources

SHRM (Society for Human Resource Management). "Publications." 2013. www.shrm.org/PUBLICATIONS/pages/default.aspx

Wells, J. T. *Corporate Fraud Handbook: Prevention and Detection*. 3rd ed. Wiley, 2011.

45.7 NOTES

1. For guidance on setting policies, see Chapter 44 in this *Handbook*; for details of email and Internet usage policies, see Chapter 48; for security awareness advice, see Chapter 49; and for applications of the principles of social psychology in reinforcing a culture of security, see Chapter 50 in this *Handbook*.
2. Peter Levine, quoted in www.virtualhrscreening.com/background/whybackground.htm (URL inactive).
3. Department of Justice Press Release, "Jury Finds Former Coke Employee Guilty in Conspiracy to Steal and sell Coca-Cola Trade Secrets," February 2, 2007, http://atlanta.fbi.gov/dojpressrel/pressrel07/tradesecrets020207.htm (URL inactive).
4. See Chapter 10 in this *Handbook* for more details on understanding computer crime statistics.
5. "Around the Nation: 2d Man Pleads Guilty in Wells Fargo Case," *The New York Times*, August 12, 1981, p A.10 (fee required), http://tinyurl.com/5byd6c

6. L. J. Kushner, "Career Management 101 for Information Security Pros," SearchSecurity Website, June 29, 2006, http://searchsecurity.techtarget.com/generic/0,295582,sid14_gci1196912,00.html
7. L. Stack, "Employees Behaving Badly: Combating Desk Rage," HC Online Website, October 5, 2004, www.hcamag.com/resources/learning-and-development/employees-behaving-badly-combating-desk-rage-110891.aspx
8. See Chapter 47 in this *Handbook*.
9. See Chapter 39 in this *Handbook*.
10. See Chapter 47 in this *Handbook*.
11. Nicholas Leeson's official Website: www.nickleeson.com/biography/index.html
12. M. E. Kabay, "Continuity of Operations and the Capability Matrix," *InfoSec Perception*, February 1, 2013, http://resources.infosecskills.com/perception/continuity-of-operations-and-the-capability-matrix (URL inactive).
13. M. E. Kabay, "Computer-Aided Consensus," 2009, www.mekabay.com/methodology/cac.pdf
14. M. Jacka, and P. Keller, *Business Process Mapping: Improving Customer Satisfaction*. Wiley, 2009.
15. S. Pacenka, "Computer Crime?" State of Oregon v. Randal Schwartz, Washington County Circuit Court C94-0322CR, 2007. Complaint brought by Mr. Schwart's client, the Intel Corporation; www.lightlink.com/spacenka/fors
16. I. Winkler, "Guard against Titan Rain Hackers," *Computerworld*, October 20, 2005, www.computerworld.com/securitytopics/security/story/0,10801,105585,00.html
17. N. Thornburgh, "The Invasion of the Chinese Cyberspies (and the Man Who Tried to Stop Them)," *Time*, August 29, 2005, www.time.com/time/magazine/printout/0,8816,1098961,00.html (URL inactive).
18. J. Vijayan, "Reverse Hacker Wins $4.3M in suit against Sandia Labs: Shawn Carpenter used his own hacking techniques to probe outside breach," *Computerworld*, February 14, 2007, www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9011283
19. M. Krausz, "Terrifying Your Employees: Not Recommended for Training," InfoSec Perception Website, March 12, 2012, http://resources.infosecskills.com/perception/terrifying-your-employees-not-recommended-for-training (URL inactive).
20. S. Terlap and E. Morath, "Final Day at Ford Bittersweet for Scores of Salaried Workers," *Detroit News*, March 1, 2007, http://detroitnews.com/apps/pbcs.dll/article?AID=/20070301/AUTO01/703010359/1148

CHAPTER 46

VULNERABILITY ASSESSMENT

Rebecca Gurley Bace and Jason Sinchak

46.1 THE SCOREKEEPER OF SECURITY MANAGEMENT 46·1

46.1.1 What Is Vulnerability Management? 46·1
46.1.2 What Is Vulnerability Assessment? 46·2
46.1.3 Where Does Vulnerability Assessment Fit in Security Management? 46·3
46.1.4 Brief History of Vulnerability Assessment 46·3

46.2 A TAXONOMY OF VULNERABILITY ASSESSMENT TECHNOLOGIES 46·4

46.2.1 Vulnerability Assessment Strategy and Techniques 46·4
46.2.2 Network Scanning 46·4
46.2.3 Vulnerability Scanning 46·5
46.2.4 Assessment Strategies 46·5
46.2.5 Strengths and Weaknesses of VASs 46·6
46.2.6 Roles for Vulnerability Assessment in System Security Management 46·7

46.3 PENETRATION TESTING 46·7

46.3.1 Testing Goals 46·8
46.3.2 Testing Perspectives 46·9
46.3.3 Testing Types 46·9
46.3.4 Social Engineering 46·10
46.3.5 Managing a Penetration Assessment 46·11

46.4 FUTURE READING 46·12

46.5 NOTES 46·13

46.1 THE SCOREKEEPER OF SECURITY MANAGEMENT. Information security has, over time, evolved from a collection of esoteric security issues and technical remedies to its current state, in which it is more tightly integrated with the area of enterprise risk management. One hallmark of this evolution from technology to management discipline is the emphasis placed on the deployment and use of vulnerability management practices. Vulnerability management has three complementary functional components, *vulnerability assessment* (VA), *penetration testing*, and *remediation*. Vulnerability management is considered fundamental to modern information security practice, and its components have adapted in architecture, features, and interfaces to accommodate the changing landscape of modern enterprises.

46.1.1 What Is Vulnerability Management? Vulnerability management is the process of assessing deployed IT systems in order to determine the security state of those systems. It includes the automated scanning of system attributes (vulnerability assessment), the manual testing and exploitation of systems in search of illicit

authorization or access (penetration testing), and the determination of corrective measures to mitigate issues identified that represent exposures for the enterprise, as well as managing the application of those measures (remediation). Vulnerability assessment is considered by many the key technology component of vulnerability management. However, there is a synergy between VA and the other elements of vulnerability management such that these elements drive what features are supported by VA technology (i.e., vulnerabilities that arise in the course of penetration testing may be reflected in future versions of VA scanners). Although early vulnerability management systems focused on vulnerabilities in operating systems and network configuration of IT enterprise systems, over time, the scope of VM coverage has grown to include applications as well.

The vulnerability management process is often defined in terms of four key functions. They are as follows:

- ***Inventory.*** Before an examiner can determine the extent of testing, she must identify all systems that are resident within the domain of interest. At this time, the operating system platforms and functions associated with each system are articulated and documented; furthermore any unauthorized or unmanaged systems are identified.
- ***Focus.*** Once an examiner has an idea of the systems that reside within the network, he must determine what information he (or his assessment tools) needs to see in order to find those vulnerabilities that are relevant to him and the enterprise. Some include the tuning of VA tools as a part of this step. When VA is driven by compliance requirements, such information is often specified in the regulations or policies in question.
- ***Assess.*** The examiner performs any testing (both automated and manual) necessary to identify vulnerabilities resident on the systems. She then assesses the results of these tests. Finally, she evaluates (and ranks) the actual risk to her organization's systems security, making that judgment guided by security policy and current risk management criteria.
- ***Respond.*** Finally, the examiner must execute procedures that act on the results of the assessment in order to address the problems identified. As much of current vulnerability management is performed as part of regulatory compliance, formal reporting of the results of VA, including remediation status, is an important part of this step.

Now that we've laid the cornerstones of the vulnerability management process, let's proceed to the technology particulars of VA.

46.1.2 What Is Vulnerability Assessment? VA is the analysis of the security state of a system on the basis of system information collected on demand. The four-step strategy for VA is as follows:

1. A predetermined set of target system attributes (for example, specific parameters for particular firewalls), are sampled.
2. The sampled information is placed in a data store.
3. The data store is organized and compared to a reference set of attributes.
4. The differences between the data store and the reference set are documented and reported.

46.1.3 Where Does Vulnerability Assessment Fit in Security Management? Vulnerability assessment and its parent function, vulnerability management are key elements of virtually all modern system security management strategies. There are many reasons for the popularity and acceptance of these functions, including:

- When systems are first deployed, VA allows you to baseline the security state of those systems.
- When security breaches are suspected, VA may allow you to quickly identify likely paths of attack and furthermore, determine whether they have been exercised.
- When new vulnerabilities are reported, VA can allow you to identify systems that are subject to the vulnerabilities so that you can patch or otherwise address the exposures.
- The results of individual VAs can be archived and used to document the security state of systems at a specific time. This is often used to satisfy regulatory audit or other oversight requirements.

Although VA and vulnerability management are accepted as important system security functions, they are not sufficient to protect systems from all security threats. Such measures should be included in a more comprehensive security strategy that includes security policy and procedural controls (see Chapters 44, 47, 50, 51, 52, and 53 in this *Handbook*), network firewalls and other perimeter defenses (Chapters 26 and 27), strong identification and authentication mechanisms (Chapters 28 and 29), access control mechanisms (Chapter 32), file and link encryption (Chapters 7, 32, and 33), file integrity checking (Chapters 7, 24, and 37), physical security measures (Chapters 22 and 23), and security training and education (Chapters 43, 45, 48, and 49). There are references and insight into building the rest of a comprehensive security strategy throughout this *Handbook*.

46.1.4 Brief History of Vulnerability Assessment. Early vulnerability assessment systems (VASs) include the COPS system, developed in the late 1980s by Eugene H. Spafford and Daniel Farmer at Purdue University.[1] COPS was a UNIX-targeted credentialed VA product that gained wide acceptance in security circles. The initial freeware version of the Internet Security Scanner (ISS) was also released in the early 1990s, as was Farmer's and Wietse Venema's VA software, SATAN.[2]

Subsequent trends of note include the advent of open-source tools for performing various forms of VA (e.g., NESSUS,[3] which provides vulnerability scanning capabilities, and NMAP,[4] which provides an inventory of the systems on a specific network, along with the network services resident on those systems).

Also of note is the move of VA from individual systems to the network infrastructure (i.e., dedicated network-connected appliances) and to the cloud (i.e., the VA is actually performed over the network, controlled by the security administrator over a Web interface, with results stored offsite).

Finally, as endpoints have become more mobile and application-centric than classic workstations, VA coverage has grown to cover software applications. The solution provider market for VA has adapted to these trends, and modern VA practice usually involves a mix of application software, dedicated appliances, managed services, and expert professional services (especially in areas such as penetration testing.)

46.2 A TAXONOMY OF VULNERABILITY ASSESSMENT TECHNOLOGIES.

VA products snapshot the security state of systems, diagnosing problems that indicate that the system is vulnerable to specific attacks.

VA performs an examination of key indicators within systems that are known to correspond to security vulnerabilities. Some consider some types of VAs a special case of a host-based, interval-based intrusion detection process and describe it in terms of the intrusion detection process model outlined in Chapter 27 in this *Handbook*.

46.2.1 Vulnerability Assessment Strategy and Techniques.

As noted above, there is a four-stage high-level strategy for conducting VA: Sample, Store, Compare, and Report. This four-stage process becomes quite complex when fleshed out to cover the full range of vulnerabilities that are commonly exploited by modern adversaries. A standard VA process may involve the use of any or all of the following techniques:

- ***Network Scanning.*** Maps network nodes and associated services by use of a "port scanner." This articulates issues at a network and network services layer of abstraction.
- ***Vulnerability Scanning.*** Takes the port scanning functions of network scanners to the next level by testing for operating system and application software system vulnerabilities resident on hosts connected to the network.
- ***Password Cracking.*** Identifies weak passwords.
- ***Log Review.*** A feature rooted in the earliest production computing platforms, log review remains one of the most powerful means of identifying weaknesses in systems.
- ***Integrity Checking.*** Uses checksums, hash totals, and digital signatures to allow quick and reliable detection of tampered files and system objects.
- ***Virus Detection.*** Scans for known viruses infecting systems.
- ***War Dialing.*** Scans enterprise systems for unauthorized modems.
- ***War Driving.*** Scans enterprise systems for unauthorized wireless LAN connections.
- ***Penetration Testing.*** Reenacts attackers' behavior in order to gain access and thereby test technical and procedural security measures.[5]

As many of these techniques are complex enough to merit a full chapter of their own (and are covered in specific chapters in this *Handbook*), we focus on only three of them in this chapter: network scanning, vulnerability scanning, and penetration testing.

46.2.2 Network Scanning.

Network scanners run a function called a *port scanner*, which uses a feature of ICMP, part of TCP/IP, in order to identify hosts available in a specified network address range. The port scanner then scans identified systems for open network ports and associated services.

Some network scanners use inference and other techniques to make intelligent guesses about the operating systems being run on open systems. These inferences are based on the combinations of active ports observed on a specific host (e.g., if Port 80 is open on a given host, it is likely running a Web server). Some scanners listen to ports for traffic that provides additional information about the system connected to that port.

One example of this type of surveillance is called *banner grabbing*, and can provide a great deal of detail about the connected system. The scanner collects information that is routinely sent to any process on connection; this *banner* often includes details of the name and version of the application.[6]

As a primary security policy decision involves specifying what systems are allowed on an enterprise network, network scanners provide a vital piece of information for VA. They also provide collateral information regarding the connected systems on a network, which is critical for tuning and refining the VA process. It is important, however, to understand that although the operation of a network scanner is automated, the identification of vulnerabilities discovered by network scanning is not. There is minimal, if any, decision support in network scanning products—some will identify certain port numbers as associated with a known Trojan attack. Therefore, it is important that the results from a network scanner be evaluated by someone familiar with the assessed network and associated security policy.

Increasingly, VA products include passive network scanning in addition to active network scanning. This involves placing passive network monitors to sniff traffic between target systems, inferring vulnerabilities from the traffic patterns observed between the target systems. This is useful in certain enterprise situations when protective mechanisms interfere with more classic active scanning (e.g., personal firewall agents on endpoints) but provides insights versus actionable results.

46.2.3 Vulnerability Scanning. Vulnerability scanning is the heart of traditional VA systems. In some ways vulnerability scanning appears to be the same as port scanning, but it differs in a critical way—it takes the additional step of not only collecting data regarding traffic, connections, and system attributes, but also analyzes the data to determine whether it matches a known vulnerability. Many systems (credentialed; see below) also attempt to correct the vulnerabilities identified in the scan, either automatically or overseen by human operators.

As vulnerability scanning usually targets specific hosts, it often conducts a deeper inspection than network scanners, identifying software versions, specific software applications, and configuration attributes of systems. The policy checks available to host-based vulnerability scanners can include usage patterns for software.

Perhaps the most valuable feature of vulnerability scanners is the current database of known vulnerabilities. Virtually all commercial offerings in this area include updates to these vulnerability databases as a core feature of the product. Many products offer features that assist security managers in configuring scanners to fit their environments, including a wide variety of configuration, reporting, and support features.

Vulnerability scanners can perform extremely fine-grained security assessment, in far more detail than network port scanners. However, this degree of detail comes at a price. Vulnerability scanners are typically slower and more resource greedy than network port scanners. The range of available assessment techniques is usually richer for vulnerability scanners; some (e.g., DDoS testing) can disrupt the normal operation of an enterprise network. False-positive rates (i.e., vulnerabilities spotted where none exist), can be high for many vulnerability scanners, requiring more human intervention. Finally as the value of the vulnerability scanner resides in its vulnerability database, it is critical that the database be updated frequently. A vulnerability scanner with a noncurrent database leaves users open to compromise.

46.2.4 Assessment Strategies. As in intrusion detection systems, VASs have features that allow differentiation between individual systems. The primary descriptors

for VASs involve the information sources and how those information sources are generated.

46.2.4.1 Credentialed Monitoring. Credentialed monitoring approaches for VAs are those that utilize system data sources such as file contents, installed applications, configuration information, and status information. This information is gained from nonintrusive sources; that is, it is gained by performing standard system status queries and inspection of system attributes. These sources are accessible only when the entity gathering them has legitimate access to the system (i.e., the entity possesses access credentials). In UNIX systems, this information is gathered at the host or device level; therefore, *credentialed approaches* are also *host-based* approaches. As many modern operating systems (e.g., Windows) handle status information differently, often providing it in the form of Application Programming Interfaces (APIs), the *credentialed = host-based* equivalence does not always apply in those environments.

46.2.4.2 Noncredentialed Monitors. Noncredentialed monitoring approaches are those approaches that stage system attacks and record target system responses to the attacks. These approaches are much more intrusive than credentialed attacks and do not assume (nor do they require) any legitimate access to the target system; they are launched from an attacker's perspective. Noncredentialed approaches are often called *active* approaches and have detection and monitoring features that complement those of credentialed approaches. In particular, noncredentialed approaches are usually superior for diagnosing vulnerabilities associated with network services. As noted, in UNIX systems, noncredentialed assessments are usually considered *network-based* assessments. For instance, network scanning is a noncredentialed monitoring process. It bears repeating that here, as in credentialed approached, the equivalence relationship does not necessarily apply in Windows and other modern operating system environments that provide status information in API form.[7]

46.2.5 Strengths and Weaknesses of VASs. Knowing that security point products (e.g., firewalls, NIDS, and access-control systems) that defend particular features of a security perimeter cannot be perfect in the dynamic world of network and system attack, VASs serve an important function in the overall strategy for protecting information assets.

The benefits associated with vulnerability analysis are as follows:

- VASs conserve time and resources, as they allow even nonexpert personnel to check systems automatically for literally thousands of problems, any one of each might result in an intrusion.
- VASs can be extremely helpful in training security novices to make systems more secure.
- VASs can be updated to reflect new knowledge of vulnerabilities found by vendors and researchers.
- VASs can be configured to address specific vulnerabilities and configurations affected by regulatory requirements. They are a critical component in IT security regulatory compliance.

As VASs are used in more environments as a part of a risk-management regime, they are helpful for benchmarking the security state of systems in order to document progress toward a protection goal.

As security problems have become an item of broad interest to customer and regulatory communities, the capabilities of VASs may be a fundamental requirement for operating in certain IT markets. This inclusion of VA as a requirement for doing business is likely to not only remain stable over time, but grow at an appreciable rate.

VASs are systematic and therefore consistent. These attributes allow them to be used as quality assurance measures for network or security managers. Many security professionals routinely recommend that operational security policies include provisions for using VASs to check systems for problems after major changes have occurred (as might be the case whenever software is updated or system recovery is required).

Weaknesses in VA include the following:

- Although VA is necessary for system security, it is not in and of itself sufficient to secure a system.
- Many VASs serve only to diagnose problems, not to correct them. They are part of a more comprehensive vulnerability management process. User follow-through is still required.
- If the VASs are not kept up to date, they may mislead users into underestimating the risk of penetration.
- VAS can negatively impact the performance of an operational network. It is critical to balance demand for the network with the performance hits associated with running some types of VAS.

As in many other knowledge-based security tools, VA can be used for either productive or malicious purposes. In the hands of a security manager, VA is a valuable diagnostic technique. In the hands of an attacker, VA may optimize efforts to identify targets of attack and provide insight as to exactly how those targets might be compromised.

46.2.6 Roles for Vulnerability Assessment in System Security Management. VA products can be used at several points in the system security management life cycle.

First, when a new program is put into place, a VA can baseline the security state of the system. This application is particularly valuable in establishing the case for a security program, as it provides hard evidence that security problems exist.

Next, when an operational system changes (as might be the case when software updates are installed or new equipment is connected), a VA can help. It can find specific security vulnerabilities that occur as side effects of such changes.

Finally, when security incidents occur or are suspected, VA results can assist investigators in diagnosing possible paths of entry for attackers, locating artifacts of attacks (such as back doors or Trojan horses), and identifying the system resources affected by an attack so that they may be restored to their original states.

46.3 PENETRATION TESTING. A penetration assessment is used to augment the vulnerability management program's reoccurring and baseline VA by iteratively exploiting known and unknown vulnerabilities with the intent of demonstrating business impact and risk across technical, physical, and personnel environments. The goal of a VA and underlying automated vulnerability scanning activities is to identify and prioritize *known* technical vulnerabilities. A penetration assessment will test the effectiveness of the VA through the objective identification of vulnerabilities in the environment and the ensuing exploitation of them. The exploitation of known vulnerabilities

demonstrates the real-life impact and business risk associated with the VA's ability to effectively identify known vulnerabilities and remediate them in the environment. The exploitation of *unknown* vulnerabilities or "stacked" known vulnerabilities will assess the organization's ability to respond to an incident or mitigate the impact caused by unforeseen vulnerabilities. Exploitation is performed in a manner that closely simulates actions that a real-life attacker is capable of performing outside and within the targeted environment.

Penetration assessments are performed by professional, ethical hackers using automated and manual techniques in conventional and unconventional ways. A penetration assessment should be leveraged to augment a VA and not as a replacement for the vulnerability scanning component. Assessments are performed by a *team* of qualified practitioners, ensuring well-rounded subject matter expertise in diverse areas.

Human penetration ("pen") testers have several advantages over purely automated methods:

- Automated VA tools, such as vulnerability scanners, are incapable of leveraging a vulnerability, component, or specification through a sequence of iterative and adaptive actions which may produce a high-impact outcome.
- A VA is only capable of testing and assessing the organization's susceptibility to known vulnerabilities; a penetration test is inclusive of unknown vulnerabilities and the exploitation of both.
- Pen testers possess an out-of-the-box mentality combined with a continuous strive to adapt circumvention efforts and pivot around preventative controls.
- Penetration tester activities are very difficult to mirror through automated means primarily due to human intuition.
- A pen tester thinks very differently than a traditional system or network administrator, making a large effort to solve problems in an *untraditional* way using *traditional* tools or methods.

46.3.1 Testing Goals

46.3.1.1 Demonstration of Impact and Risk. As penetration testers, our hearts thrive on the ability to compromise a target and received the coveted command level access. Although this may appear successful in the short term, the ultimate goal of a penetration test is to determine, demonstrate, and explain the risk as it applies to the organization, with accompanying steps to mitigate it or adequately manage it in the future.

Management should discuss with the penetration team what the assessment's end goal should be. *Goal-oriented* penetration testing aims to maximize value by demonstrating impact and risk appropriately for the organization's unique business. Goals include things that are important to the organization outside of security controls, such as the theft of intellectual property, theft of client base information, or a reduction in availability for core business processes thereby affecting an organization's bottom line.

46.3.1.2 Attestation and Compliance. The result of the penetration assessment can be provided to management, business partners, or regulators as an attestation to organizations point-in-time level of susceptibility to impact as a result of the current effectiveness of the vulnerability management program and underlying assessment capabilities.

A penetration assessment is a point-in-time assessment. Due to the fluid nature of information security advancement, a penetration testing team can never provide any level of assurance relating to the audit of vulnerabilities currently present in the environment. The primary goal of a penetration assessment is to demonstrate business impact and risk to management and business partners for the threats they face at a specific point in time.

46.3.2 Testing Perspectives. Penetration assessments provide a vehicle for evaluating the security of technical, physical, personnel, and procedural controls at a holistic level through the objective perspective of a third party mimicking the actions of an external attacker or malicious internal user. A VA is typically managed by an internal office and is therefore subject to bias during the assessment phase and compounded by resistance to the advancement in skillsets required to align with an advancing security field.

A malicious user and external attacker are different perspectives and approaches to the assessment which result in different perceived levels of impact and risk. It is important to understand the difference in order to properly develop the scope of the assessment to cover VA and management's objective for the assessment.

External Attacker

- An external attacker mimics the capabilities of an outsider with no internal knowledge of the target environment.
- The capabilities of an external attacker are greatly limited to publicly available resources.
- This type of test typically involves attacks against perimeter network devices and applications, which may result in internal connectivity.

Malicious User

- A malicious user mimics the capabilities of an authenticated or authorized user who possesses knowledge of internal operations.
- The capabilities of a malicious user are only afforded to trusted employees or contractors inside the network perimeter or within an application's authenticated or authorized zone.
- This type of test typically poses a greater risk of impact than an external attack, primarily due to the intricate knowledge pertaining to the location of and adequacy of preventive controls around sensitive data.

46.3.3 Testing Types. Penetration testing can be utilized to expand the scope of a technical VA to include the physical, personnel, and procedural controls that may be targeted by an adversary in an attack. Additionally, penetration tests can be used to increase focus on a particularly weak technical area within the over-arching assessment.

External Assessment: Penetration efforts originate from the Internet and target Internet-facing devices with the primary goal of gaining internal access to systems and acquiring business trophies such as client information or intellectual property.

Application Assessment: Application level assessments are increasingly common as organizations continue to remove Internet-accessible external services and heavily leverage Web services. Application assessments involve user-based attacks such as manipulating user input or client-controlled parameters. Goals include accessing administrative panels, extracting sensitive database information, masquerading as authorized users, or compromising the underlying Web or backend database server.

Internal Assessment: Internal penetration assessments target vulnerabilities within the organization. This assessment assumes the role of an attacker or malicious users who has been provided access to a network port within the facility. Goals of this assessment mirror those of an external or application assessment with the inclusion of a reduced level of accessibility controls.

Wireless Assessment: Wireless penetration assessments involve internally and externally surveying the organization's physical facilities for authorized and unauthorized wireless communications (e.g., 802.11) and attempting to circumvent or exploit controls to gain access. The primary goal of this assessment is to gain internal access to resources. Secondary goals of a wireless assessment include those of an internal assessment on achieving connectivity.

Physical Assessment: Physical security assessments focus testing activities on bypassing controls designed to prevent access to the facility and access to physical electronic data such as computing areas or telecommunication closet.

Personnel and Social Engineering: Personnel and social engineering assessments aim to test an organization's security awareness training and ability of personnel to follow procedures designed to safeguard their interactions with external parties.

46.3.4 Social Engineering. Security depends on people; therefore, people are the single weakest link in any penetration assessment and overarching VA. People possess the power to assist an attacker in covertly bypassing any security control in the environment with a minimal amount of resources. Social engineering is defined as the use of persuasive techniques to abuse an inherent trust relationship and induce personnel to divulge information or perform an action used to gain illicit access to systems or data. Social engineering attacks can be performed through multiple mediums of communication, such as telephone conversations, convincing emails, or spoofed Websites.

Social engineering attacks have historically been leveraged in many high-profile security breaches involving Advanced Persistent Threats (APTs). This is primarily due to the level of effectiveness, covert nature, high impact of return, and ease of execution. Including social engineering as a component of a penetration assessment should be considered to evaluate policies, procedures, and security awareness programs designed to protect the business and inform personal of this particular threat.

Practical recommendations

- The decision to include Social Engineering as part of the assessment should align with objectives and goals.
- Precautionary measures should be taken to protect the testing team from unnecessary harm if they are identified; this can be provided through a "get-out-of-jail-free" document provided by upper management.

Common assessment scenarios
The following are example scenarios commonly executed during penetration assessments inclusive of Social Engineering tests:

- *USB/CD Drops*: The testing team stages a series of USB thumb drives or CDs with an auto-run or convincing "click me" executable. This executable "calls home" to an Internet rendezvous point once executed. This test is benign and rudimentary but a valuable determination of the susceptibility of authorized personnel to executing attacker provided code on internal systems.
- *Email Phishing/Pharming*: Hoaxes or scams are sent to a subset of targeted personal groups by an attacker masquerading as an internal employee or department with the intent of coercing them into responding, visiting a malicious Web page, or executing an attachment.
- *Call Masquerading*: The testing team contacts a provided list of phone numbers for a subset of personal groups within the organization to elicit information or perform an action by posing as an employee or contractor.

46.3.5 Managing a Penetration Assessment. Varying testing models can used to fine-tune the impact of the penetration assessment and align it with the vulnerability management and assessment goals.

46.3.5.1 Announced versus Unannounced Testing. Announced testing introduces awareness to internal IT staff prior to the beginning of the assessment. The primary motivation behind announced testing is to reduce any perceived impact from testing activities.

- Due to heightened awareness, announced testing reduces the effectiveness of the assessment to accurately evaluate security monitoring and incident response processes.
- Announced testing can provide educational value by presenting a learning opportunity to internal security staff through the real-time observation of controls, which are attacked or circumvented.

Unannounced testing involves the covert performance of testing activities under the sole knowledge and authorization from upper management to perform the agreed-upon assessment. Unannounced testing will deliver the most accurate assessment of all tested security controls.

Due to the covert nature of unannounced testing, activities are commonly focused on demonstrating impact through only exploiting vulnerabilities, which will quickly and covertly lead to a desired trophy.

46.3.5.2 Testing Scope. The penetration testing scope should be developed so that it aligns with the objectives of the assessment and its purpose within the vulnerability management program. It should be noted that regulatory requirements may dictate the scope to include specific network segments or applications (e.g., the Payment Card Industry Data Security Standard, or PCI DSS).

Penetration assessments can be scoped to allow social engineering attacks and specify to what degree these attacks will proceed, such as targets for phishing campaigns or locations for USB/CD drive drops.

Scope should ensure the assessment covers all intended areas within the VA, such as physical, personnel, and procedural controls.

- Scope should be reduced to a manageable number of targets within the environment for the allocated time period to ensure proper assessment of critical components.
- Scope reduction may eliminate a subset of vulnerabilities but may be necessary to ensure adequate time is available to test the environment and important assets using the provided methodology.

46.3.5.3 Testing Methodology. A penetration assessment should be performed under a consistent and repeatable methodology to ensure it can be operated in a repeatable manner for each VA.

Penetration testing activities involve a large variety of techniques and tools. The best way to harness and document the range of activities is through a vetted and documented methodology.

A structured methodology is crucial to accurately document the scope and depth for which testing activities should be performed.

- The assessment must be performed under a proven and structured methodology to be accepted by many third-party auditors or regulatory agencies.
- A structured methodology enables the ability to transition testing resources into and out of the assessment and provides results in a comparative manner between testing periods.

46.4 FUTURE READING

Abraham, J. "Goal-Oriented Penetration Testing—The New Process for Penetration Testing." Personal Website, November 16, 2009. http://spl0it.wordpress.com/2009/11/16/goal-oriented-pentesting-the-new-process-for-penetration-testing

Abraham, J. "Goal-Oriented Penetration Testing (Part 2)." Personal Website, November 17, 2009. http://spl0it.wordpress.com/2009/11/17/goal-oriented-pentesting—the-new-process-for-penetration-testing-part-2/

Foster, J. C., and V. T. Liu. *Writing Security Tools and Exploits*. Syngress, 2006.

Fischer, M., and Kabay, M. E. "Penetration Testing, Part 3," Network World Security Strategies, February 11, 2003, www.networkworld.com/newsletters/sec/2003/0210sec1.html

Faircloth, J. *Penetration Tester's Open Source Toolkit*, 3rd ed. Syngress, 2011.

Hurley, C., R. Rogers, F. Thornton, and B. Baker. *Wardriving & Wireless Penetration Testing*. Syngress, 2007.

Kabay, M. E. "Social Engineering in Penetration Testing: Cases." *Network World Security Strategies Alert*, October 25, 2007. www.networkworld.com/newsletters/2007/1022sec2.html

Manzuik, S., A. Gold, and C. Gatford. *Network Security Assessment: From Vulnerability to Patch*. Syngress, 2006.

McGraw, G., "Is Penetration Testing a Good Idea?" Network Magazine, July 2005, www.cigital.com/papers/download/0507sec.penetration.pdf

Orlando, J., and Kabay, M. E., "Social Engineering in Penetration Testing: Analysis," Network World Security Strategies, October 30, 2007, www.networkworld.com/newsletters/sec/2007/1029sec1.html

PCI Security Standards Council. *Information Supplement: Requirement 11.3 Penetration Testing*. 2008. https://www.pcisecuritystandards.org/pdfs/infosupp_11_3_penetration_testing.pdf

Schumacher, P., and Kabay, M. E. "Social Engineering in Penetration Testing: Intimidation," Network World Security Strategies, November 8, 2007, www.networkworld.com/newsletters/sec/2007/1105sec2.html

Siemens Insight Consulting, "Penetration Testing," April 2006, www.insight.co.uk/files/datasheets/ Penetration%20Testing%20(Datasheet).pdf

Skoudis, E. "Maximizing Value in Pen Testing." *SANS Penetration Testing*, February 9, 2012. http://pen-testing.sans.org/blog/2012/02/09/maximizing-value-in-pen-testing (URL inactive).

Van Der Walt, C., H. D. Moore, R. Temmingh, H. Meer, J. Long, C. Hurley, and J. Foster. *Penetration Tester's Open Source Toolkit*. Norwell, MA: Syngress, 2005.

46.5 NOTES

1. D. Farmer and E. H. Spafford, "The COPS Security Checker System," *Proceedings of the Summer USENIX Conference,* Anaheim, California, June 1990, pp. 165–170, http://docs.lib.purdue.edu/cgi/viewcontent.cgi?article=1844&context=cstech
2. D. Farmer and W. Venema, "Improving the Security of Your Site by Breaking into It," Internet white paper, 1993, www.csm.ornl.gov/~dunigan/cracking.html
3. Nessus Vulnerability Scanner. www.nessus.org
4. NMAP.ORG. http://nmap.org/
5. J. Wack, M. Tracy, and M. Souppaya, "Guideline on Network Security Testing," NIST Special Publication 800-42, National Institute of Standards and Technology, October, 2003, http://csrc.nist.gov/publications/nistpubs/800-42/NIST-SP800-42.pdf (URL inactive).
6. wiseGEEK, "What is Banner Grabbing?" wiseGeek Website, 2013, www.wisegeek.com/what-is-banner-grabbing.htm
7. A. Shostack and S. Blake, "Towards a Taxonomy of Network Security Assessment Techniques." *Proceedings of 1999 Black Hat Briefings,* Las Vegas, NV, July 1999, www.blackhat.com/presentations/bh-usa-99/AdamS/shostack-blackhat.pdf

CHAPTER 47

OPERATIONS SECURITY AND PRODUCTION CONTROLS

M. E. Kabay, Don Holden, and Myles Walsh

47.1 INTRODUCTION 47·2
- 47.1.1 What Are Production Systems? 47·2
- 47.1.2 What Are Operations? 47·3
- 47.1.3 What Are Computer Programs? 47·3
- 47.1.4 What Are Procedures? 47·4
- 47.1.5 What Are Data Files? 47·4

47.2 OPERATIONS MANAGEMENT 47·4
- 47.2.1 Separation of Duties 47·4
- 47.2.2 Security Officer or Security Administrator 47·5
- 47.2.3 Limit Access to Operations Center 47·5
- 47.2.4 Change-Control Procedures from the Operations Perspective 47·7
- 47.2.5 Using Externally Supplied Software 47·9
- 47.2.6 Quality Control versus Quality Assurance 47·10

47.3 PROVIDING A TRUSTED OPERATING SYSTEM 47·13
- 47.3.1 Creating Known-Good Boot Medium 47·13
- 47.3.2 Installing a New Version of the Operating System 47·13
- 47.3.3 Patching the Operating System 47·13

47.4 PROTECTION OF DATA 47·14
- 47.4.1 Access to Production Programs and Control Data 47·14
- 47.4.2 Separating Production, Development, and Test Data 47·15
- 47.4.3 Controlling User Access to Files and Databases 47·15

47.5 DATA VALIDATION 47·16
- 47.5.1 Edit Checks 47·16
- 47.5.2 Check Digits and Log Files 47·17
- 47.5.3 Handling External Data 47·18

47.6 CLOUD COMPUTING AND PRODUCTION SYSTEMS 47·18

47.7 CONCLUDING REMARKS 47·20

47.8 FURTHER READING 47·21

47.9 NOTES 47·21

47.1 INTRODUCTION. Despite the enormous increase in individual computing on personal computers and workstations in the years since the first edition of this *Handbook* was published in 1975, centralized arrays of dozens, hundreds, or thousands of computers used in computational arrays and cloud services are still used for enterprise computing in applications devoted to the core business of the enterprise. This chapter focuses on how to run vital computers and networks safely and effectively.

Readers with a military background will note that operations security in the civilian sector is different from the OPSEC (a military acronym for "operational security") designation used in military discourse. As defined by the Joint Chiefs of Staff of the United States military, "OPSEC seeks to deny real information to an adversary, and prevent correct deduction of friendly plans."[1]

In determining what operations security and production controls are required for any system, a thorough risk analysis should be conducted; you will find references throughout this *Handbook.* For a quick analysis, it may be helpful to use a few other common military acronyms and list the threats from their perspectives.

EMPCOA—enemy's most probable course of action. What is the most likely course of action an attacker will take against your systems?

EMDCOA—enemy's most dangerous course of action. What is the worst possible thing an attacker could accomplish?

Weighing these as part of a risk analysis can help tremendously to decide how to employ limited resources. The acronym *METT-TC*, which is essentially a larger version of the engineering triad of cost, time, and quality, may help. There are various versions of the METT-TC acronym, but perhaps the most useful stands for

- Mission (what we need to do),
- Equipment (what we have),
- Time (by when we need to do it),
- Troops (whom we have to do it),
- Terrain (where are we doing it), and
- Culture (the possible cultural considerations).

Each part of our METT-TC analysis should be self-evident, except for perhaps culture. Ignoring the management, political, and community cultures of the locations in which to install security controls is a frequent mistake, usually with dire consequences for both security and the careers of security professionals.

It is critical to be able to define precisely the terms associated with operations security and production controls before engaging in a discourse about them.

47.1.1 What Are Production Systems? A *production system* is one on which an enterprise depends for critically important functions. Examples include systems for handling accounts receivable, accounts payable, payroll, inventory, manufacturing systems, real-time process control, data-entry systems, Web-based client interfaces for e-commerce, critical information systems, portable data-handling systems, and management information systems. What is defined as critical is a function of the mission, not a uniform prescription. Thus, the accounts receivable may be a critically important function for a manufacturing company working on a tight profit

margin, with receivables needed to pay immediately for crucial supplies, but it may be less critical for a small group of successful partners running a consulting firm.

47.1.2 What Are Operations? *Operations* consist of the requirements for control, maintenance, and support of production systems. Operations staff are responsible for such functions as:

- Integrating new software systems into an existing configuration
- Running programs and batch jobs to update databases and create reports
- Installing new versions of production programs
- Maintaining production databases for maximal efficiency
- Managing backups (creation, labeling, storage, and disposal)
- Responding to emergencies and recovering functionality
- Mounting storage volumes of tapes, cartridges, or disks in response to user or program requests
- Handling special forms for particular printouts (e.g., check blanks)
- Managing all aspects of production networks, such as configuring routers, bridges, gateways, wireless propagation, and firewalls

47.1.3 What Are Computer Programs? A *computer program* is a set of instructions that tells a computer what to do to perform a task. Computer programs may be acquired, or they may be internally developed.

Internally developed programs are stored in computer systems in two basic forms. *Source* programs are in the form in which they were written (coded) by computer programmers. The statements in source programs are in languages such as COBOL, Visual BASIC, C++, and Java. Source language programs are kept in files and stored on disk folders called *source libraries* or *program libraries*.

Executable programs have been converted from source code, by *compilation* or *interpretation*, into a program that the computer can execute. Executable programs may be maintained in two separate forms: *object* and *load*.

An object program is a partially executable module that must be linked to other executable modules, such as input/output modules, to become a load module. As load modules, the programs are said to be in executable form. Executable programs are kept in *production libraries*, from which they are called when needed. Acquired programs are generally in object and load form. The source code is proprietary to the organization that developed it and is rarely given to the acquiring enterprise.

When internally developed programs have to be changed, programmers work with copies of the source programs. The copies are stored in another type of library, referred to as *programmer libraries*. The programmers make changes to their copy of the source programs, and go through a process of recompiling and testing until the modified program is working properly. When acquired programs require changes, often a contract to make the modifications is issued to the organization from which the programs were acquired. In some situations, internal programmers generate new programs and interfaces to the original acquired programs. The same libraries are used: source libraries for source code, production libraries for executable modules, and programmer libraries for work in progress. These libraries need to be protected. Loss or damage can entail huge costs and considerable inconvenience; recovering them can require a long time and great expense.

47.1.4 What Are Procedures? *Procedures* are sets of statements that tell a computer what to do in certain situations. They are unlike programs in that they are not compiled. Stored in files or databases, they are invoked and interpreted as needed. Procedural statements are made up of operational commands and parameters. The operational commands tell the computer what to do, and the parameters tell the computer which entity to act upon. *Job Control Language* (JCL) is an example of a procedural language. Procedural language statements often are used in database management systems and in security software products. On personal computers, *batch files* (.BAT) are a form of simple procedure.

47.1.5 What Are Data Files? Almost everything stored and maintained in a computer system takes the form of a *file*.[2] Programs, procedures, information, all are stored in files, using the concept of a file in its broadest sense; that is, a collection of related items. This usage has become most apparent with the ubiquitous personal computer (PC). Data files, as distinguished from program files and other types, are those that store information. In a PC environment, data files may be called documents. Documents are created by word processors, spreadsheet programs, graphics generators, and other application programs. In mainframe and midsize computer environments, data files are those created and maintained by applications such as payroll, accounting, inventory, order entry, and sales.

Some data files are transient; that is, they are created, used, and deleted within a short period of time. If lost or damaged, they can be reconstructed quickly, with little difficulty. There is usually no need to protect transient files. Other files, such as *master files* or organizational *databases* (groups of files that are linked to one another using a database management system or DBMS), contain information that is vital, confidential, or virtually irreplaceable. These files, generated by PCs, mainframes, and midsize computer systems, must be protected by security software and backup procedures to ensure against loss, destruction, theft, interference, and unauthorized disclosure.

47.2 OPERATIONS MANAGEMENT. The processes for effective and efficient management of operations have direct benefits on information assurance. In particular, these aspects of operations management are of special value for improving and maintaining security:

- Separation of duties
- Defining the role of the security officer or security administrator
- Limiting access to the operations center
- Defining secure change-control processes
- Careful controls over externally supplied software
- Managing quality assurance and quality control

47.2.1 Separation of Duties. *Separation of duties* (also discussed in Chapter 45 in this *Handbook*) is a key control that should be applied to development and modification of programs. In enterprises where there are systems and programming departments that create and maintain custom programs, each individual programmer is assigned a user ID and a password. In these enterprises, where programs are developed and maintained internally, changes are constantly made to programs in order to meet

changing business requirements. Modified executable programs, after recompilation and testing by programmers, are moved from their libraries into production libraries. Modified source programs are moved into source libraries. The programmers are responsible for keeping the source libraries current, while computer operations, or some other functional group separated from programming, may be responsible for maintaining the production libraries. When an updated executable program is transferred from the programmer's library to the production library, a transmittal is included, signed off by a manager in the programming department.

A particular consequence of the separation of duties is that a member of the operations staff should always be involved in the functional analysis and requirements definition phases for changes to production programs. The operations perspective is not always clear to programmers, and such issues as logging, backout, and recovery, as discussed later in this chapter, need to be brought to their attention early in the development and maintenance cycles.

It is critically important that programmers and operations staff understand that under no circumstances are programmers to install changes into production systems without adequate testing and authorization. Patch management is discussed in detail in Chapter 40 in this *Handbook*.

47.2.2 Security Officer or Security Administrator. Contemporary enterprises typically include a mix of external cloud-computing services, mainframes, midsize computers, and local area networks (LANs) comprising hundreds or thousands of workstations, PCs, terminals, and other devices, all interconnected with one another, and often connected with the same mix in other enterprises throughout the world via the Internet.

A department, or an individual or small group in smaller enterprises, has the responsibility for providing and maintaining the security of files, databases, and programs. The title that is often associated with this function is *information security officer* (ISO) or *information systems security officer* (ISSO). This individual or department has the mandate to carry out the security policy as set down by the senior management of the enterprise. The security officer is empowered to allow or to disallow access to files, databases, and programs. In the language of the security officer, procedures are set up and maintained that establish relationships among individuals, programs, and files. Users, programmers, and technicians are granted privileges for full access, update only, or even read only. The security officer has the power to change or to revoke privileges (see Chapter 24 in this *Handbook* for details of operating-system security mechanisms; see Chapters 28 and 29 for discussions of identification and authentication).

47.2.3 Limit Access to Operations Center. Physical access to the operations center grants a person enormous power to disrupt production systems. Such access must be tightly controlled.

47.2.3.1 Need, Not Status, Determines Access. A fundamental principle for effective security is that access to restricted areas is granted on the basis of roles. Employees whose roles do not justify access should be excluded from autonomous access to production systems. In particular, high-placed executives, such as the president, chief executive officer, chief financial officer, chief operating officer, chief technical officer, and all vice presidents, should examine their own roles and determine if they should be able to enter the operations center unaccompanied; in most cases, such access is unjustified. Limiting their own access sets an important model for other aspects

of security policy and demonstrates that need, not social status or position within the corporate hierarchy, determines access to restricted areas. This issue is also discussed in Chapter 45 on employment practices and policies in this *Handbook*.

47.2.3.2 Basic Methods of Access Control. As explained in Chapters 28 and 29 in this *Handbook*, access control depends on identification and authentication (I&A). I&A can be based on:

- What one has that others lack (tokens such as physical keys or smart cards)
- What one knows that others don't know (user IDs and passwords or passphrases)
- What one is that differs from what others are (static biometric attributes, such as fingerprints, iris patterns, retinal patterns, and facial features)
- What one does differently from how others do it (dynamic biometrics, such as voice patterns, typing patterns, and signature dynamics)

A typical arrangement for secure access to an operations center may involve keypads for entry of a particular code or card readers programmed to admit the holders of specific magnetic-stripe cards, smart cards, or radio-frequency identification (RFID) tokens. If the operations center is a 24-hour operation with full-time staffing, the presence of operators provides an additional layer of security to preclude unauthorized access. Remote monitoring of sensitive areas increases security by discouraging unauthorized access or unauthorized behavior and speeds up the response to possible sabotage. For extensive discussion of physical and facilities security, see Chapters 22 and 23 in this *Handbook*.

47.2.3.3 Log In and Badge Visitors. Visitors to a facility that houses sensitive systems should be logged in (that is, their identification checked and recorded in writing) at a controlled entrance and provided with visitor badges.

In high-security applications, an additional login may be required when entering the operations center itself. To encourage return of visitor badges, some security policies require the visitor to deposit a valuable document, such as a driver's license, with the security guards at the main entrance.

The effectiveness of visitor badges as a means of identifying nonemployees depends entirely on the use of badges by all personnel at all times; if *not* wearing a badge is acceptable and common, a malicious visitor could simply hide a visitor badge to pass as an authorized employee.

The time of login and of logout can be valuable forensic evidence if malfeasance is detected. However, such records can be shown to be reliable only if the guards responsible for keeping the logs consistently verify the completeness and correctness of all information written into the logs. Video recordings of the entrance to capture the face of each visitor as well as the exact time of arrival and departure can be of great value both for identifying imposters and also for verifying that the security guards are doing their job properly.

47.2.3.4 Accompany Visitors. No unaccompanied visitors should be permitted to circulate in the operations center or in the facility housing such a center. In high-security facilities, someone must even accompany the visitor to the washroom and wait for the visitor outside the door.

If a consultant or temporary employee is to work on a project for longer than a day, it may be acceptable to grant that person a restricted pass for low-security areas; however, high-security areas, such as the operations center, would still require such a person to be accompanied.

When one of the authors (Kabay) was director of technical services for a large service bureau in Montreal in the mid-1980s, the president of the company arrived unannounced on a Saturday evening at the door of the operations center with a visitor to show him around the computer center. The young operator explained with some trepidation that the president's name was not on the list of approved solo visitors, and so he would not grant entry (although he offered to call his boss for authorization). To his credit, the president accepted the restriction gracefully—and wrote a letter of commendation for the operator on Monday morning.

47.2.4 Change-Control Procedures from the Operations Perspective.

When programmers have made changes to production programs and all documentation and testing procedures are complete, the new versions are formally turned over to the operations staff for integration into production.

47.2.4.1 Moving New Versions of Software into Production.

Operations managers and staff must meet these demands when moving new versions of software into production:

- Identification—tracking which software is in use
- Authorization—controlling changes
- Scheduling—minimizing disruptions to production
- Backups—ensuring that all requisite information is available to restore a prior state
- Logging—keeping track of data input for recovery, and of errors for diagnosis of problems
- Backout—returning to a prior production version in case of catastrophic errors

47.2.4.1.1 Identification. Knowing precisely which versions of all production software are in use is the basis of production controls. Every module must have a unique identification that allows immediate tracking between executable code and source code; all changes to a particular module must be fully documented by the programming group. Unique identifiers allow the quality assurance process to ensure that the only modules that go into production are those that have been properly tested.

Most production shops use a three-level numbering scheme to track versions. Typically, version *a.b.c* (e.g., 7.13.201) is defined in this way:

- *c* changes every time anything at all—even a spelling mistake—is changed.
- *b* changes when program managers decide to group a number of fixes to errors into a new version for release to production.
- *a* changes when significant new functions are added; often the source code is completely renumbered if the changes are great enough.

The version number of object code must match the number of its source code. All object code should include internal documentation of its version number so that

the version can be ascertained instantly, without having to consult possibly inaccurate external documentation.

47.2.4.1.2 Authorization. Strict procedures must be in place to preclude rogue programmers from introducing modified code into the production suite. In addition to the dangers of introducing untested or undocumented changes, allowing any individual to modify production processes without verification and authorization by an appropriate chain of responsibility can allow hostile code, such as Trojan horses and backdoors (see Chapters 13, 16, and 20 in this *Handbook*), to be introduced into the systems.

47.2.4.1.3 Scheduling. Implementing any new version of a production system requires careful planning and scheduling. Operations staff must prepare for changes in all aspects of production that depend on the system in question; for example, there may be requirements for new printer forms, additional magnetic tapes, and other supplies that must be ordered in advance. New requests for operator intervention, or changes in status and error messages during production, necessitate appropriate documentation and training. Effects on other programs may require special preparations that must take into account the scheduling requirements of the other systems that are affected. In addition, new versions of software often are implemented immediately after major production jobs, such as end-of-year or quarterly processing, to maintain consistency within an accounting period. For all these reasons, scheduling is critically important for trouble-free operations.

47.2.4.1.4 Backups. When modifying production systems, operations staffs usually take one or more complete backups (see Chapter 57 in this Handbook) of the software and data to be modified. This procedure is essential to allow complete restoration of the previous working environment should there be catastrophic failure of the new software and data structures.

System managers and operators should be aware that if their proposed changes require reformatting disk drives on the production system, they must make at least *two* full backups to ensure that an error in one of the backups won't cause a catastrophe when the entire system has to be restored.

47.2.4.1.5 Logging. To allow recovery or backout without losing the new data and changes to existing data that may have been carried out using new software and data structures, all production programs should include a logging facility. Logging keeps a journal of all information required to track changes in data and to regenerate a valid version of the data by applying all changes to an initial starting condition. Logging requires synchronization with backups to avoid data loss or data corruption. Special requirements may exist when a new version of the production system involves changes to data structures; in such cases, applying the information about changes to the older data structures may require special-purpose application programs. Since programmers sometimes forget about such possibilities, operations staff should be prepared to remind the programming staff about such requirements during the design phases for all changes.

47.2.4.1.6 Backout and Recovery. Sometimes a new version of production software is unacceptable and must be removed from the production environment. This decision may be made immediately, or it may occur after a significant amount of data entry and data manipulation has taken place. In either case, operations should be able

to return to the previous version of a production system without data loss. This process involves restoring the earlier complete operating environment, with software and data in synchrony, and then using log files to repeat the data input and data modifications from the moment of changeover to the moment of fallback (see Chapter 52 in this *Handbook*).

Not all of the changes that were made using a new version will necessarily be applicable to the previous data; for example, if new fields were added to a database, the data stored in those fields would not be usable for an older, simpler data structure. Similarly, if fields were removed in the newer database, recovery will involve providing values for those fields in the older database. All of these functions must be available in the recovery programs and backups that should accompany any version change in production systems.

47.2.4.2 Using Digital Signatures to Validate Production Programs. If an unauthorized intruder or a disgruntled employee were discovered to have gained access to the production libraries, it would be necessary to determine if there had been unauthorized modifications to the production programs.

Date and time stamps on programs can record the timing of changes, but many operating environments allow such information to be modified using system utilities that read and write directly to disk without passing through normal system calls. In those instances, there would be no time stamp or log entry.

One approach that has been used successfully is to apply checksums to all production components. Checksum software applies computations to programs as if the codes were simply numbers; the results can be sensitive to changes as small as a single bit. However, if the checksums are computed the same way for all programs, access to the checksum utility could allow a malefactor to change a module and then run the checksum utility to create the appropriate new checksum, thus concealing the evidence of change. To make such subterfuge harder, the checksums can be stored in a database. Naturally, this database of checksums itself must be protected against unauthorized changes, for example, using encryption. Storing checksums may make unauthorized changes more difficult to disguise, but it also extends the chain of vulnerabilities.

A better way of determining whether object code or source code has been modified is to use digital signatures. Digital signatures are similar to checksums, but they require input of a *private key* that can, and must, be protected against disclosure. Verifying the digital signature may be done using a corresponding *public key* that can be made available without compromising the secrecy of the private key. For more information on public and private keys, see Chapter 7 on encryption and Chapter 37 on PKI and certificate authorities in this *Handbook.*

When digital signatures are used to authenticate code, it may be possible to validate production systems routinely, provided that the process is not too arduous to be accomplished as part of the normal production process. For example, it should be possible to validate all digital signatures in no more than a few minutes, before allowing the daily production cycle to start.

47.2.5 Using Externally Supplied Software. Production often uses software from outside the organization; such software may be commercial off-the-shelf (COTS) programs or it may consist of programs modified for, or written especially for, the organization by a software supplier. In any case, external software poses special problems of trust for the production team. There have been documented cases in which production versions of software from reputable software houses have contained viruses

or Easter eggs (undocumented features, such as the well-known flight simulator in MS-Excel 97, which popped up a graphic landscape that included a monitor showing the names of the Excel development team). In addition, some consultants have publicly admitted that they deliberately include Trojan horse code (undocumented malicious programming) that allows them to damage data or inactivate the programs they have installed at client sites if their fees are not paid.

In large data centers, it may be possible to run quality-assurance procedures on externally supplied code (see Chapter 39 in this *Handbook*). Such tests should include coverage monitoring, in which a test suite exercises all the compiled code corresponding to every line of source code. However, it is rare that an operations group has the resources necessary for such testing.

The trustworthiness of proprietary external software written or adapted especially for a client ultimately may depend on the legal contracts between supplier and user. Such legal constraints may not prevent a disgruntled or dishonest employee in the supplier organization from including harmful code, but at least they may offer a basis for compensation should there be trouble.

47.2.5.1 Verify Digital Signatures on Source Code If Possible. If externally supplied code is provided with its source library as well as with compiled modules, operations should try to have the supplier provide digital signatures for all such programs. Digital signatures will permit authentication of the code's origins and may make it harder for malefactors to supply modified code to the user. In addition, the digital signatures can support nonrepudiation of the code (i.e., the supplier will be unable credibly to claim that it did not supply the code) and therefore the signatures may be useful in legal action, if necessary.

47.2.5.2 Compile from Source When Possible. Wherever possible, it is highly desirable to be able to compile executables from source code on the target machine. Compiling from source allows quality assurance processes to check the source for undocumented features that might be security violations and to couple the executables tightly to the verified source. In addition, compiling on the local system ensures that all calls to system routines will be satisfied by linking to executables such as dynamic link libraries (DLLs) supplied in the current version of the operating system.

However, compilation on a local system has additional implications that complicate implementation of new applications: Because the system routines being linked to the compiled code may not be identical to those used during the manufacturer's quality assurance tests, the customer organization must plan for its own quality assurance testing.

Operations staff should express this preference for source code clearly to the person or group controlling acquisition of external software. For more information about writing secure code, see Chapter 38 in this *Handbook*; for information about secure software development and software quality assurance, see Chapter 39.

47.2.6 Quality Control versus Quality Assurance. Throughout this chapter, quality assurance has been mentioned as an essential underpinning for operations security. *Quality assurance* refers to the processes designed to ensure and to verify the validity of production programs. However, another aspect of quality concerns the operations group: the quality of output. The process of verifying and ensuring the quality of output is known as *quality control.*

47.2.6.1 Service-Level Agreements. Unlike mathematical truth, there is no absolute standard of quality for computing operations. Every organization must define the level of quality that is suitable for a particular application. A commonly quoted principle in programming and operations is that there is a complex relationship among quality, cost, and development time: Increasing quality increases both cost and development time; shortening development time increases cost, if quality is to be maintained. It follows that every system should include a definition of acceptable performance; such definitions are known as *service-level agreements* (SLAs).

SLAs typically include minimum and maximum limits for performance, resource utilization, and output quality. The limits should be expressed in statistical terms; for example, "The response time measured as the time between pressing ENTER and seeing the completed response appear on screen shall be less than three seconds in 95 percent of all transactions and shall not exceed four seconds at any time." SLAs may define different standards for different types of transactions if the business needs of the users so dictate.

47.2.6.2 Monitoring Performance. Computer-system performance depends on five elements:

1. Access time and speed of the central processing unit(s) (CPU)
2. Access time and speed of mass storage (disks)
3. Access time and speed of fast memory (RAM)
4. Application design
5. Available network bandwidth

Operations groups should monitor performance to ensure that the requirements of the SLAs are met. There are two approaches to such monitoring: (1) analysis of log files and (2) real-time data capture and analysis.

Log files that are designed with performance analysis in mind can capture the precise times of any events of interest; for example, one might have a record in the log file to show when a particular user initiated a read request for specific data and another record to show when the data were displayed on the user's screen. Such level of detail is invaluable for performance analysis because the data permit analysts to look at any kind of transaction and compute statistics about the distribution of response times. In turn, these data may be used for trend analysis that sometimes can highlight problems in program or data structure, design, or maintenance. For example, an excessively long lookup time in a data table may indicate that the system is using serial data access because it had been designed without an appropriate index that would permit rapid random access to the needed records.

Another approach to performance monitoring and analysis is to use real-time monitors that can alert operations staff to abnormal performance. For example, an application program may be designed to calculate response times on the fly; the results may be displayed numerically or graphically on a *dashboard* for the operations staff. Values falling below a specified parameter may signal an abnormal condition, using color or sound to alert the operators to the drop in performance. Such *integrated performance metrics* allow the fastest possible response to performance problems. Extensive sets of real-time measures are sometimes called *situational awareness* tools.

Even if the application programs lack integrated performance metrics, it is sometimes possible to use system-level online performance tools to analyze system activity. In one instance in the experience of one of the authors (Kabay), for example, a software supplier in the mid-1980s had promised a response time of 10 seconds or less for all transactions, but one particular operation was taking 43 minutes. Because the SLA was violated, the client was threatening to sue the supplier for the entire development cost and the three years of maintenance payments—a total of $3.8 million (about $12 million in 2013 dollars). Using an online performance tool, it quickly became obvious that the transaction in question was generating an enormous amount of disk I/O (read and write operations): 80,000 random-access reads in a particular data set to filter out a few target records using a nonindex field. Installing and using an appropriate index and compacting the data set on that field as the primary key to provide rapid access to blocks of related records reduced response time to 6 seconds.

47.2.6.3 Monitoring Resources. Consistent monitoring of resource utilization is one of the most valuable roles of the operations staff. Data center operations should include regular analysis of system log files to track changes in the number of files, amount of disk free space available, number of CPU cycles consumed, number of virtual memory swap operations, and less esoteric resource demands, such as numbers of lines or pages printed, number of tape mounts requested, number of backup tapes in use, and so on. These data should be graphed and subjected to trend analysis to project when particular resources will be saturated if the trend continues. Even a simple spreadsheet program such as Excel can produce regression lines in graphs easily. Operations can then reduce demand either by improving aspects of production (e.g., optimizing programs to require fewer resources) or by increasing available resources (e.g., installing a memory upgrade).

Another level of analysis focuses on specific users and groups of users. Each functional group using the production systems should be analyzed separately to see if there are discontinuities in their trends of resource utilization. For example, a specific department might show a relatively slow and stable rise in CPU cycles consumed per month—until the rate of increase suddenly increases tenfold. If such a rate of increase were to continue, it could surpass all the rest of the system demands combined; operations therefore would investigate the situation before it caused problems. The cause of the discontinuity might be a programming error; for example, there might be a logical loop in one of the programs or a repeated computation that ought to have its result stored for reuse. However, the change in slope in CPU utilization might be due to introduction of new programs with a rapidly growing database; in such cases, operations would have to act to meet heavy new demands.

Disk space is often a key resource that can cause problems. If users fail to clean up unwanted files, disk space can disappear at an astounding rate. This problem is exacerbated by poor programming practices that allow temporary work files to remain permanently in place. Systems have been designed with tens of thousands of sequentially numbered temporary work files that had no function whatsoever after a production run was completed but that were accumulated over several years.

One of the methods widely used to reduce resource waste is *chargeback.* Using system and application log files, system administration charges the users of particular systems a fee based on their use of various resources. Sometimes these chargebacks are viewed as *funny money* because they are an accounting fiction—no money actually changes hands. However, requiring managers to budget carefully for computing resources can greatly improve attention to mundane housekeeping matters such as

cleaning up useless files. If the chargeback system extends to aspects of program performance such as number of disk I/Os, it can even influence programmers to optimize their design and their code appropriately. Optimization is appropriate when the total costs of optimization are less than the savings in resources and the increases in productivity that result from optimization efforts, measured over the lifetime of the application.

47.2.6.4 Monitoring Output Quality. The final component of quality control is the meticulous monitoring of everything that is produced in the data center and sent to users or clients. Although much printing is now performed on local printers controlled by users, in many situations the operations group is responsible for documents such as payroll checks, invoices, and account status reports. Every operations group must explicitly assign responsibility for verifying the quality of such output before it leaves the data center. Operators should keep careful logs that record various types of error (e.g., torn paper, misaligned forms, or poor print quality) so that management can identify areas requiring explicit attention or repairs to improve quality.

47.3 PROVIDING A TRUSTED OPERATING SYSTEM.

The operating system (OS) is usually the single biggest and most important example of externally supplied software in a data center. Because the OS affects everything that is done in production, it is essential to know that the software is trustworthy. To this effect, operations staff use procedures to ensure that *known-good* software is always available to reinstall on the system. Systems with such software are known as *trustworthy computing platforms*.

47.3.1 Creating Known-Good Boot Medium.

The simple principle that underlies known-good operating software is that there shall be an unbroken chain of copies of the OS that have never run any other software. That is, operations will create a boot medium (tape, cartridge, CD-ROM) immediately after installing known-good software.

For example, if boot-medium V1B0 is defined as version 1 of the OS as it is delivered from the manufacturer; its installation would require specific settings and parameters for the particular configuration of the system. Immediately after installing V1B0, but before running any other software, operations would create medium V1B1 and set it aside for later use if V1B0 had to be replaced.

47.3.2 Installing a New Version of the Operating System.

Continuing this example of how to maintain known-good operating software, it might become necessary to install a new version of the operating system—say, version 2 on medium V2B0. Before using V2B0, operations would reinstall the current known-good OS, say from V1B1. Only then would V2B0 be installed, and the new boot medium V2B1 would be created immediately.

47.3.3 Patching the Operating System.

Often, when it is necessary to modify a small part of the OS, rather than installing a whole new version, manufacturers ask users to *patch* the OS. The patch programs modify the compiled code in place. If checksums or digital signatures are in use to maintain OS integrity, these codes will

have to be regenerated after the patch is applied. However, to maintain a known-good status, applying a patch should follow a rigid sequence:

1. Load the current known-good software from the appropriate medium (e.g., V2B1).
2. Install the patch.
3. Immediately create a known-good boot medium before running any other software (in our example, this medium would be V2B2).

For extensive discussion of managing patches for production systems, see Chapter 40 in this *Handbook*.

47.4 PROTECTION OF DATA

47.4.1 Access to Production Programs and Control Data. Just as the operations center needs restricted access, so do production programs and data. From a functional point of view, there are three categories of people who might be allowed access to programs and data on which the enterprise depends: users, programmers, and operations staff.

47.4.1.1 Users. The only people who should have read and write access to production *data* are those users assigned to the particular systems who have been granted specific access privileges. For example, normally only the human resources staff would have access to personnel records; only the finance department staff would have full access to all accounts payable and accounts receivable records. Managers and other executive users would have access to particular subsets of data, such as productivity records or budget figures. Of course, *no* user should ever have write access to production *programs.*

47.4.1.2 Programming Staff. Programmers create and maintain production programs; they naturally have to be able to access the versions of those programs on which they currently are working. However, programmers must not be able to modify the programs currently used in *production*. All changes to production programs must be documented, tested, and integrated into the production environment with the supervision of quality assurance, operations, and security personnel.

Programmers need to be able to use realistic data in their development, maintenance, and testing functions; however, programmers should *not* have privileged access to restricted data. For example, programmers should not be allowed to read confidential files from personnel or medical records or to modify production data in the accounts payable system. Programmers can use extracts from the production databases, but particular fields may have to be randomized to prevent breaches of confidentiality. Programmers generally resent such constraints, but usually an effective process of education can convince them that maintaining barriers between production systems and systems under development is a wise policy.

47.4.1.3 Operations Staff. Much as the programmers are responsible for developing and maintaining systems, so the operations staff are responsible for using and controlling these systems. Operations staff perform tasks such as scheduling, error handling, quality control, backups, recovery, and version management. However,

operations staff should not be able to modify production programs or to access sensitive data in production databases.

47.4.2 Separating Production, Development, and Test Data. For obvious reasons, testing with production data and production programs is an unacceptable practice, except in emergency situations. Therefore, programmers who develop new programs or modify existing programs must perform tests using their own libraries. These are frequently referred to as test libraries. Experienced programmers keep copies of the source programs for which they are responsible as well as copies of some of the data files that are used by the programs and subsets of others in their own test libraries. To avoid security violations, such copies and subsets of privileged data should be anonymized to the degree necessary to protect confidentiality. For example, a system using personnel data might substitute random numbers and strings for the employee identifiers, names, and addresses.

It is important to include time stamps on all files and programs, including both production and test versions. This practice serves to resolve problems that arise about program malfunctions. If all programs and files have time stamps, it can be helpful in determining whether the most current version of the load program is in the production library and whether test files and production files have been synchronized.

Final testing prior to production release may entail more formal review by an independent quality assurance section or department; the quality assurance group may also control transfers of programs to the production library.

47.4.3 Controlling User Access to Files and Databases. Access to files has to be controlled for two reasons:

1. There is confidential information in files that is not to be made accessible to everyone.
2. There are other files that are considered auditable.

The information in these files may be confidential, but that is not the reason for controlling access to them. The information in these files must be controlled because changing it is illegal. An example would be an enterprise's general ledger file once the books have been closed. Changing the information in these files gave birth to the pejorative phrase *cooking the books*. The original copies of these files are developed on the computer that handles the day-to-day transactions of an enterprise. Some time after the month-end closing of the books, copies of these files are archived. In this form they are the recorded history of the enterprise and cannot be changed. Storing this chiseled-in-stone historical information, combining details and summaries in database format so as to be accessible for the purpose of analysis is known as data warehousing.

In most large enterprises, these files are created on mainframe and midsize computers, although microcomputer database servers in LANs are increasingly in use today. In any event, controlling user access to files is performed in several ways depending on what types of access are allowed. Remote access to online production databases and files is often done over leased lines—dedicated communication facilities paid for on a monthly basis as opposed to dial-up or switched lines paid for on a when-used basis. Access control may be readily accomplished through the use of front-end security software modules, which in turn feed into database and file handling software and finally into the application software. For example, in an environment using a software product

for handling queries and updates of online databases and files, a number of different security software products could be installed. Such products use what are called rules or schemas to validate user IDs and passwords, to authorize types of transactions, and to allow access to files and databases.

Many information system installations allow remote communications for many kinds of transactions. Various individuals, including sales representatives entering order information and traveling executives wishing to access current information from databases or files, use public networks, such as wired or wireless Internet service providers at hotels, airports, and coffee shops. This type of access increases the potential for security breaches. The usual practice today is to use *virtual private networks* (VPNs), which consist of encrypted channels between the portable equipment and the enterprise networks. Even so, some inadequate implementations of VPNs allow cleartext transmission of the initial logon information, allowing the identification and authentication data to be captured and used for unauthorized access. Poorly secured network access also allows *man-in-the-middle* attacks on the user's traffic.

For more information about encryption, see Chapters 7 and 37 in this *Handbook*; for more about network and communications security, see Chapters 5, 25, 32, 33, and 34.

47.5 DATA VALIDATION. Just as it is essential to have trusted operating systems and application software for production, the operations group must be able to demonstrate that data used for production are valid.

Validation controls normally are carried out dynamically throughout data entry and other processing tasks. Some validity checks are carried out automatically by database software; for example, inconsistencies between header records and details may be reported as errors by the database subsystems. Bad pointers are usually flagged immediately as errors by the database software; examples include:

- Pointers from a particular master record to a detail record with the wrong key value or to a nonexistent location
- Forward or backward pointers from a detail record to records that have the wrong key value for the chain or that do not exist at all

However, many errors cannot be caught by database subsystems because they involve specific constraints particular to the application rather than errors in the database itself. For example, it may be improper to allow two chemical substances to be mixed in a processing vat, yet there is nothing in the data themselves that the database software would recognize as precluding those two values to be recorded in the input variables. The programmers must include such restrictions in edit checks; often these relations among variables can be coded in a data dictionary. If the programming environment does not allow such dependencies, the programmers must incorporate the restrictions in lookup tables or in initialization of variables.

From a production point of view, operations staff must run appropriate validation programs created by the database suppliers and by the application programmers to assure the quality of all production data. The next sections review in more detail what is involved in such validation programs.

47.5.1 Edit Checks. Operations should have access to diagnostic programs that scan entire databases looking for violations of edit criteria. For example, if a field is

designated as requiring only alphanumeric characters but not special characters such as "#" and "@," then part of the diagnostic sweep should be checking every occurrence of the field for compliance with those rules. Similarly, range checks (greater than, less than, greater than or equal, equal, less than or equal, between) are a normal part of such scans. Lookup tables listing allowed and forbidden data and combinations of data provide further sophistication for more complex relations and restrictions. In any case, the role of operations staff is to run the diagnostics and identify errors; correction of the errors should fall to authorized personnel, such as the database administrators.

Diagnostic programs should provide detailed information about every error located in the production files. Such details include:

- Configurable view of the record or records constituting an error, showing some or all of the fields
- Unique identification of such records by file name or number, record number, and optionally by physical location (cylinder, sector) on disk
- Error code and optional full-text descriptions of the error, including exactly which constraints have been violated

A diagnostic program should, ideally, also allow for repair of the error. Such repair could be automatic, as, for example, insertion of the correct total in an order-header, or manual, by providing for the database administrator to correct a detail record known to be wrong.

47.5.2 Check Digits and Log Files. Another form of verification relies on check digits. Programs can add the numerical or alphanumeric results of data manipulations to each record or to groups of records when transactions are completed properly. Finding records with the wrong check digits will signal inconsistencies and potential errors in the processing. Check digits are particularly useful to identify changes in production databases and other files that have been accomplished through utilities that bypass the constraints of application programs. For example, most databases come with a relatively simple *ad hoc* query tool that permits lookups, serial searches, views, and simple reporting. However, such tools often include the power to modify records in compliance with database subsystem constraints, but completely free of application program constraints.

An even more powerful type of utility bypasses the file system entirely, and works by issuing commands directly to the low-level drivers or to the firmware responsible for memory and disk I/O. In the hands of the wrong people, both database and system utilities can damage data integrity. However, it is usually difficult for the users of these utilities to compute the correct checksums to hide evidence of their modifications. A diagnostic routine that recomputes checksums and compares the new values with the stored values can spot such unauthorized data manipulations immediately.

If possible, checksums in log files should be *chained* from one record to the next so that the checksum for every record is calculated using the checksum from the previous record in addition to the data in the current record. Changing, adding, or deleting any record in the log file then requires the criminal to modify all subsequent checksums, and comparison of the modified log file with the backup of the original log file may easily identify the tampering because there are so many discrepancies in evidence.

A similar technique for validating and repairing data uses database and application log files to record information such as before and after images of modified records.

Such log files also can include special marker records to flag different steps in complex transactions; these flags can allow diagnostic programs to identify precisely when transactions have been interrupted or subjected to other forms of failure. Using these log files, it is often possible to identify which defective transactions need to be removed, repeated, or completed.

Commercial data-integrity software is available for a wide range of platforms and databases. Searching on "data integrity software" in the Google search engine (www.google.com) locates many references to such products.[3]

For more information on integrating security into application design, see Chapter 52 in this *Handbook*. For more about log files and other monitoring and control systems, see Chapter 53.

47.5.3 Handling External Data. Before using data provided by external organizations, operations should routinely check for data purity. Diagnostic routines from the programming group should be available to check on all data before they are used in batch processing to update a production database. The same principles of data validation used in checking production databases should apply to all data received from clients, suppliers, governments, and any other organization. Special validation programs can and should be written or obtained to test the data received on any medium, including tapes, cartridges, removable discs, CD-ROMs, DVDs, and data communications channels.

47.6 CLOUD COMPUTING AND PRODUCTION SYSTEMS. The history of production computing is repeating itself: In the 1960s through the 1980s, computers were so physically large and hugely expensive that many smaller organizations contracted with *service bureaus* to access computing resources. Connections in physically close locations (e.g., city cores) were often through physical coaxial cable or twisted pair connections. For more distant connections, clients linked their *dumb terminals* to the *mainframes* through telephone lines using *modems.* Sometimes the switched telephone connections through the *plain old telephone service* (POTS) were fixed in place and dedicated to the connection—hence they were called *dedicated lines.*

Today, as we approach the middle of the 2010s, the speed of Internet connections is reaching 10 Gbps—an enormous bandwidth that facilitates remote access to banks of computing power of almost unimaginable power.[4] Organizations are capitalizing on the possibility of creating virtual machines (*virtualization*) that insulate concurrent processes from each other, allowing far more efficient sharing of centralized resources than running processes on dedicated systems in-house. Inexpensive computers (*thin clients*) with relatively little application software and local disk storage can be used to access all the necessary programs and data required for the organization's business through access to cloud services.

In March 2013, industry analysts said, "More than 60% of all enterprises will have adopted some form of cloud computing in 2013, according to Gartner Research. The cloud market is slated to grow to $131 billion worldwide this year, or 18.5% over the $111 billion last year. In its 2013 *State of the Cloud Report* that surveyed over 1200 IT professionals, IT-reseller CDW found 39% of organizations are already using some form of cloud solution, an 11% increase over 2011."[5]

IDC reported that "In coming years, the economic impact of cloud computing will be vast. Each year a greater percentage of businesses IT budgets are earmarked for cloud. An expansive study by the International Data Corporation (IDC) … reported that, in 2011, businesses spent $28 billion on public cloud IT services. Amazingly, the

spending on public cloud expected to surpass $207 billion worldwide by 2016 … ." The researchers analyzed trends in specific industries:

- Banking: increasing use of cloud computing
- Healthcare: slower adoption
- Manufacturing: particularly strong growth for customer relationship management (CRM) and among smaller businesses
- Insurance: increasing use
- Communications/media: particularly strong user of *storage-on-demand*[6]

Gartner also predicted that, "… by 2015, 10% of overall IT security enterprise capabilities will be delivered in the cloud, with the focus today clearly on messaging, Web security, and remote vulnerability assessment. However, there's also the expectation there will be more on the way, such as data-loss prevention, encryption, authentication available too as technologies aimed to support cloud computing mature."[7]

As with service bureaus of yesteryear, cloud computing poses special challenges for operations security and production controls.

- The reliability of employees hired to handle confidential and critically important data is out of the hands of the client organization.
- Management policies, monitoring, software maintenance, audits—all are potentially handled exclusively by employees of the cloud-computing provider.
- Bring-your-own-device (BYOD) practices are facilitated by access to remote cloud services.
- Quality of service (QOS) issues and details of the service-level agreements (SLAs) complicate the contractual relations between providers and customers.

Cloud computing, like service bureaus, can provide cost-effective growth paths for smaller business and can offload information technology used for IT functions that are not viewed as mission critical, allowing IT staff to concentrate on innovative, highly productive applications that can differentiate an organization in its marketplace. They allow for graded increases in computing power without forcing organizations to follow step-functions with large investments in much bigger equipment and increased operational costs. However, extending one's IT infrastructure into centers run by separate entities with their own profit motives requires careful attention to security. At a minimum, organizations should implement the following recommendations for maintaining adequate controls over their production environment when it is in a remote site using cloud computing:

1. During evaluations of multiple vendors, be sure to contact clients of each firm to have personal discussions of their experience with the providers' service-level agreements and performance, openness of communications about production controls, cooperation in site visits, and adequacy and cooperation in resolving problems.
2. Examine the adequacy of encryption for all stored proprietary data. No cleartext data should be accessible to employees of the cloud-hosting company at any time—including while decrypted on the virtual machines running on their servers.

3. Be sure that virtual private networks are in place for all Internet-based data transfers.
4. Explicitly discuss update standards for all the software your organization plans to run on the cloud service. Are you responsible for such updates or is the cloud vendor?
5. Be sure that all software running in the cloud on behalf of the customer organization respects the terms of the vendors' licenses. For example, be sure that a one-user license is not being applied to a thousand concurrent virtual machines on your behalf.
6. Understand and analyze the business-continuity planning (BCP) and disaster-recovery planning (DRP) in place to ensure continued operations on your behalf should there be problems at the cloud vendor's site(s). Are the exact terms spelled out to your satisfaction in the contracts? Are there provisions for testing the adequacy of the BCP and DRP?
7. Ensure that the contract allows for external audits which can be initiated by the client. Independent evaluation of the security, QOS and continuity of operations (CoO) is essential for the protection of the client.
8. Discuss the vendors' security policies and practices, including real-time monitoring for breaches (*situational awareness*), handling malware, and vulnerability analysis, including penetration testing.
9. Evaluate the billing processes carefully: what determines the periodic invoicing—concurrent users? Total number of sessions? Detailed algorithms measuring such elements as disk I/O, CPU cycles, swapping to and from virtual memory, or bandwidth utilization? Model the costs if possible using detailed information from your own existing systems.

47.7 CONCLUDING REMARKS. Up until the mid-1980s, the world of mainframes and minicomputers differed from that of PCs. However, from the mid-1980s and into the millennium, these worlds have merged. LANs have proliferated, the Internet has changed the way business is conducted, and bridges, routers, and gateways make it possible for information to move among computer platforms, regardless of type. Security requirements are now universal in scope. Numerous layers of software and hardware separate the user and the technician from the information to which they require access. These layers themselves contribute to security because they require some technical skill to get at the information. As this is being written in 2008, computer literacy is increasing rapidly. Children are taught to use computers in grade school, and tens of millions of workers routinely use PCs or workstations daily, so the security provided by the technology is not as significant as it once was.

Security is also an economic issue. If an individual with the requisite skill is determined to gain access to online files or databases, it is extremely expensive to prevent such access. Even with high expenditures, success in achieving complete security is never guaranteed. Nevertheless, if the value of the information and its confidentiality justifies additional expense, there are software products available that employ complex schemes to support security. When necessary, information security can be extended down to the field level in records of online files and databases.

Other related security measures, such as physical protection, communication security, encryption of data, auditing techniques, system application controls, and other topics, are covered in other chapters of this *Handbook*. No one measure can stand alone

or provide the ultimate protection for security, but with a proper balance of measures, the exposures can be contained and managed.

47.8 FURTHER READING

Baumann, W. J., J.T. Fritsch, and K. J. Dooley. *Network Maturity Model: An Integrated Process Framework for the Management, Development and Operation of High Quality Computer Networks*. Outskirts Press, 2007.

Benyon, R., and R. Johnston *Service Agreements: A Management Guide*. Van Haren Publishing, 2006.

Blanding, S. *Enterprise Operations Management Handbook*, 2nd ed. Auerbach, 1999.

Cisco. "Cisco on Cisco Best Practices: Data Center Operations Management." White paper, 2008. www.cisco.com/web/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Cisco_IT_Operational_Overview_Data_Center_Management.pdf

Erl, T., R. Puttini, and Z. Mahmood. *Cloud Computing: Concepts, Technology & Architecture*. Prentice-Hall, 2013.

Franklin, C., and B. Chee. *Securing the Cloud: Security Strategies for the Ubiquitous Data Center*. Auerbach, 2013.

Halpert, B. *Auditing Cloud Computing: A Security and Privacy Guide*. Wiley, 2011.

Hoesing, M. T. *Virtualization Security Audit and Assessment*. Auerbach, 2014.

Mather, T., S. Kumaraswamy, and S. Latif. *Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance*. O'Reilly Media, 2009.

McCrie, R. *Security Operations Management*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 2006.

Nielsen, L. *The Little Book of Cloud Computing Security, 2013 Edition*. New Street Communications, 2013.

Rhoton, J., J. De Clercq, and D. Graves. *Cloud Computing Protected: Security Assessment Handbook*. Recursive, 2013.

Rosado, D. G., D. Mellado, E. Fernandez-Medina, and M. Piattini, eds. *Security Engineering for Cloud Computing: Approaches and Tools*. IGI Global, 2012

Winkler, J. R. *Securing the Cloud: Cloud Computer Security Techniques and Tactics*. Syngress, 2011.

47.9 NOTES

1. United States Joint Chiefs of Staff, "Information Operations," Joint Publication 3–13, 2006, p. II-2, www.dtic.mil/doctrine/jel/new_pubs/jp3_13.pdf (URL inactive).
2. There are some forms of data that don't reside in files, but they are transient. For example, memory buffers, stacks, and registers contain data used during execution of programs, but they are not normally accessible to users except through special tools such as *privileged-mode debug* utilities or special forensic programs.
3. The site http://dmoz.org/Computers/Software/Databases/Data_Warehousing/Data_Integrity_and_Cleansing_Tools listed 38 such tools at the time of writing (May 2013).
4. L. Bandoim, "Growth of Cloud Computing and ERP Continues to Accelerate," *Technorati | Technology | Cloud Computing*, May 4, 2013, http://technorati.com/technology/cloud-computing/article/growth-of-cloud-computing-and-erp
5. A. Nain, "With The Cloud Market Set To Flirt With 20% Growth In 2013, How Can You Play It?" *Seeking Alpha*, March 20, 2013, http://seekingalpha.com/article/

1291461-with-the-cloud-market-set-to-flirt-with-20-growth-in-2013-how-can-you-play-it

6. J. Weeks, "Vertical Markets—2013 Growth Predictions for Cloud Computing," *US Signal Blog*, January 16, 2013, http://ussignalcom.com/blog/vertical-markets-2013-growth-predictions-for-cloud-computing
7. E. Messmer, "Gartner: Growth in Cloud Computing To Shape 2013 Security Trends: Gartner Predicts by 2015, 10% of Overall IT Security Enterprise Capabilities Will Be Delivered in the Cloud," *Network World*, December 6, 2012, www.networkworld.com/news/2012/120612-gartner-cloud-security-264873.html

CHAPTER 48

EMAIL AND INTERNET USE POLICIES

M. E. Kabay and Nicholas Takacs

48.1 INTRODUCTION **48·2**

48.2 DAMAGING THE REPUTATION OF THE ENTERPRISE **48·2**

48.2.1 Violating Laws 48·3
48.2.2 III-Advised Email 48·3
48.2.3 Inappropriate Use of Corporate Identifiers 48·4
48.2.4 Blogs, Personal Websites, and Social Networking Sites 48·5
48.2.5 Disseminating and Using Incorrect Information 48·5
48.2.6 Hoaxes 48·6

48.3 THREATS TO PEOPLE AND SYSTEMS **48·12**

48.3.1 Threats of Physical Harm 48·12
48.3.2 Pedophiles Online 48·13
48.3.3 Viruses and Other Malicious Code 48·13
48.3.4 Spyware and Adware 48·14

48.4 THREATS TO PRODUCTIVITY **48·15**

48.4.1 Inefficient Use of Corporate Email 48·15
48.4.2 Mail Storms 48·22
48.4.3 Buying on the Web 48·24
48.4.4 Online Gambling 48·26
48.4.5 Internet Addiction 48·28
48.4.6 Online Dating and Cybersex 48·28
48.4.7 Games and Virtual Reality 48·29
48.4.8 Changing Email Addresses 48·30

48.5 LEGAL LIABILITY **48·31**

48.5.1 Libel 48·31
48.5.2 Stolen Software, Music, and Videos 48·31
48.5.3 Plagiarism 48·32
48.5.4 Criminal Hacking and Hacktivism 48·33
48.5.5 Creating a Hostile Work Environment 48·33
48.5.6 Archiving Email 48·36

48.6 RECOMMENDATIONS **48·36**

48.6.1 Protecting Children 48·37
48.6.2 Threats 48·37
48.6.3 Hate Sites 48·38
48.6.4 Pornography 48·38
48.6.5 Internet Addiction 48·38
48.6.6 Online Dating 48·39
48.6.7 Online Games 48·40
48.6.8 Online Purchases 48·40
48.6.9 Online Auctions 48·41
48.6.10 Online Gambling 48·41
48.6.11 Preventing Malware Infections 48·42
48.6.12 Guarding against Spyware 48·42
48.6.13 Junk Email 48·43

48.6.14	Mail Storms	48·43	**48.7**	**CONCLUDING REMARKS**	**48·45**
48.6.15	Detecting Hoaxes	48·44			
48.6.16	Get-Rich-Quick Schemes	48·44	**48.8**	**FURTHER READING**	**48·45**
48.6.17	Hacking	48·45	**48.9**	**NOTES**	**48·46**

48.1 INTRODUCTION.[1] The Internet offers every enterprise exciting opportunities to find timely information and to reach potential clients. This very power brings with it risks of damaging corporate and professional reputations. Nontechnical problems in cyberspace include bad information, fraud, loss of productivity, and violations of civil and criminal law as well as violations of the conventions of proper behavior established by custom in cyberspace.

In addition, widespread abuse of Internet access while at work is forcing recognition that clear policies are essential to guide employees in appropriate use of these corporate resources. The consensus in our profession—despite the dreadful lack of hard statistics—is that something like two-thirds of all the damage caused to our information systems is from insiders who are poorly trained, careless, or malicious. (For a detailed discussion of security statistics, see Chapter 10 in this *Handbook.*) For example, a study published in late 2005 reported that:

> Sixty-nine percent of 110 senior executives at Fortune 1,000 companies say they are "very concerned" about insider network attacks or data theft, according to a study by Caymas Systems, a network security technology firm based in San Jose, Calif. And 25 percent say they are so concerned they can't sleep at night, Sanjay Uppal, a vice president at Caymas Systems, told eSecurityPlanet.[2]

A McAfee-sponsored survey in Europe showed that (in the words of the Department of Homeland Security *Daily Open Source Infrastructure Report*[3]):

> Workers across Europe are continuing to place their own companies at risk from information security attacks. This "threat from within" is undermining the investments organizations make to defend against security threats, according to a study by security firm McAfee. The survey, conducted by ICM Research, produced evidence of both ignorance and negligence over the use of company IT resources. One in five workers let family and friends use company laptops and PCs to access the Internet. More than half connect their own devices or gadgets to their work PC and a quarter of these do so every day. Around 60 percent admit to storing personal content on their work PC. One in ten confessed to downloading content at work they shouldn't. Most errant workers put their firms at risk through either complacency or ignorance, but a small minority are believed to be actively seeking to damage the company from within. Five percent of those questioned say they have accessed areas of their IT system they shouldn't have while a very small number admitted to stealing information from company servers.[4]

Another topic of growing significance is saturation by floods of email sent by well-intentioned employees who do not know how to use email effectively.

Finally, some of the information in this chapter may help security administrators involve their users in a more active role by giving them take-home messages that can help them protect their own families and friends. Getting employees to care about security for their families is a good step to involving them in corporate security.

For more information on effective security awareness and corporate culture change, see Chapters 49 and 50 in this *Handbook*.

48.2 DAMAGING THE REPUTATION OF THE ENTERPRISE. When someone posts information to the 'Net, the message header normally indicates who the

sender is. In particular, all employees using a corporate email account identify their employer in every posting. It follows that when an employee—for example, joe@acme.com—misbehaves on the 'Net, it is likely that everyone seeing the misbehavior will associate it with the employer, regardless of futile attempts to dissociate the employee from the employer by statements such as "The opinions above are not necessarily those of my employer."

Employees can embarrass their employers by using their corporate email identifiers in these ways:

- **Flaming.** Launching rude verbal attacks on others.
- **Spamming.** Sending junk email (spam), unsolicited advertising, and sales promotions, to multiple, often unrelated, Usenet groups and mailing lists and to people's email addresses without their permission.
- **Mail-bombing.** Sending many email messages to a single email address to annoy its user, or in extreme cases, to cause a denial of service.

In addition, employees can violate laws, send out embarrassing content via email, implicate their employers in personal affairs, and spread falsehoods with actionable consequences.

48.2.1 Violating Laws. Employees may engage in illegal activities that can seriously compromise their employer; examples include:

- Industrial espionage
- Stock manipulation
- Criminal hacking, unauthorized penetration of other systems
- Sabotage, denial of service attacks
- Vandalism, defacement of Websites
- Creating, transmitting, or storing child pornography
- Sending threats (e.g., of harm to the President of the United States)
- Credit card fraud, using stolen or fraudulently generated credit card numbers for purchases made using corporate resources

Corporate Internet usage policies should explicitly forbid any of these actions.

48.2.2 Ill-Advised Email. There have been too many cases of foolish use of email in recent years. Employees have created a hostile working environment by sending internal email with lewd or hateful jokes and images; staff members have insulted their bosses or their employees in email that later became public; people have made libelous accusations about other workers or about competing companies. All of these uses are wholly inappropriate for a medium that takes control of distribution away from the originators and produces records that can be backed up and archived for indefinite periods of possible retrieval.

Common sense dictates that anything sent via email should not be illegal or even embarrassing if it were published in a newspaper.

Users should also be aware that it is a poor idea to insult people using email. Sending *flames* that belittle, ridicule, and demean other people is likely to generate more of the

same in response, and flaming is an ugly practice that distorts standards for public and private discourse. If a user or employee chooses to respond to a rude or demeaning email, he or she should refrain from replying with the same rude or demeaning tone. In one case reported in a security management class to author Kabay in 2013, an employee noted that a young intern responded to an email memorandum by hitting REPLY ALL and sending a contemptuous, arrogant, demeaning, and ungrammatical criticism of the message. That message had been sent to all employees by the CEO. The intern was fired that day.

Employees should work to maintain the moral high ground by refraining from obscenity, profanity, and vulgarity in written as well as in oral discourse. Not only is this a good habit in general, but it also avoids the possibility of enraging total strangers who may be physically or electronically dangerous.

These best practices also apply to the home and family life. Criminal hackers have been known to damage credit ratings, participate in identity theft to rack up large bills in the victims' names, and even tamper with phone company accounts. In one notorious prank, hackers forwarded all incoming phone calls for the famous security expert Donn Parker, who is quite bald, to a hair restoration business.

Anonymizers are services that strip identifying information from email and then forward the text to the indicated targets. However, even anonymizers respond to subpoenas demanding the identity of people involved in libel or threats. The Website called *annoy.com* consistently posts messages that will annoy a substantial number of people as an exercise of United States First Amendment rights; however, even that service once had a particularly clear message on its refusal to tolerate abuse:

> WARNING
>
> It has come to our attention that certain people have been using annoy.com to deliver what some might consider to be threats of physical violence or harm to others.
>
> Do not mistake our commitment to freedom of speech for a license to abuse our service in this manner.
>
> We plan to cooperate fully with law enforcement agencies in whatever efforts they make to find you and punish you—even if it's some renegade authoritarian dictatorship … Free speech and annoy.com are not about harassment and definitely not about harm or violence. If you think for a second we will allow cowardly idiots to spoil our free speech party you are making a mistake. A huge mistake.

For both USENET, a global Internet discussion system, and for discussion groups on the Web, a message may be forever. There are archives of USENET messages stretching back for decades, and the Wayback Machine (named for a time machine run by Mr. Peabody [a scholarly dog] in Hanna-Barbara's *Rocky* [a flying squirrel] *and Bullwinkle* [an engaging moose] TV cartoon show of the 1960s) for the Web has records back to 1996. Sending abusive or degrading messages online may not permanently damage the sender's reputation, but it is not likely to *improve* anyone's prospects for getting or keeping a good job, especially if the sender's email address includes a corporate affiliation.

48.2.3 Inappropriate Use of Corporate Identifiers. Considerable controversy exists as to whether corporate policy should forbid corporate IDs for any personal use on the Internet. There is little reason for posting messages to newsgroups in the .alt hierarchy, and especially not to groups catering to or sympathetic to criminal activity. If employees of an organization want to participate in vigorous

political discussion, conversations about sexual activity, and any other topic unrelated to their work, they are free to do so using their own Internet identities. Employers pay for corporate email identities; people who want to post opinions—especially political opinions—about, say, basket-weaving techniques should pay for their own access and leave their employer out of the postings.

The risks of damaging an organization's reputation by violating netiquette are high. Some abusers have themselves been abused by angry and unscrupulous Internauts. In one notorious early case, back in 1994, a naïve executive spammed the 'Net—he posted messages in a couple of dozen newsgroups. In retaliation, his company's 800-number was posted to phone-sex discussion groups in the .alt hierarchy, resulting in thousands of irate and expensive phone calls by seekers of aural sex. Regular customers were unable to get through, and some staff resigned because of the offensive calls. The executive nearly lost his job.

An additional risk is that employees will inadvertently post company-confidential information to what they erroneously perceive as closed, private groups. Competitors or troublemakers can then exploit the information for competitive advantage or publicize it to harm the enterprise. Even if a discussion group or mailing really is closed, nothing prevents a participant from using or disseminating confidential information without permission. By the time the breach of security is discovered, it can be too late for remediation. Such breaches are an invitation for spear-phishing attacks, as discussed in Chapter 20 in this *Handbook*.

48.2.4 Blogs, Personal Websites, and Social Networking Sites. Should employers be concerned about the creation of blogs, personal Websites, and social-networking pages by employees? There have been cases in which employees made unwise or frankly derogatory comments about their current employers, with predictable consequences. It is much better to prevent such conflicts by establishing clear policies for employees that explicitly ban mention of the employer's name in personal publications and media such as blogs and Websites. A variation can require corporate approval by the public relations or communications departments before material is published. Such policies are commonplace for control over what employees publish in interviews, newsletters, and other publications.

Some employees likely have personal pages in social-networking sites such as Facebook, Twitter, and Tumblr. The same issues arise when employees make reference to their employer by name as part of their profile: How would an employer feel about seeing an indecently dressed pictured on their Facebook page with the individual's profile displaying their corporate name? Employment agreements can, and should, stipulate limitations on the use of corporate identity. There is nothing wrong with stipulating that social-networking pages not include the name of an employer.

Further complicating the matter is the growth of professional social networking sites such as LinkedIn. These sites encourage posting a *digital resume*, which includes basic information about an individual's employment history. Much like personal networking sites, companies must set clear expectations on use and periodically monitor compliance through searches and visual inspection. In contrast with Facebook, though, LinkedIn tends not to have many users who post unprofessional information and comments on their own pages.

48.2.5 Disseminating and Using Incorrect Information. The Internet and in particular the World Wide Web are in some ways as great a change in information distribution as the invention of writing 6,000 years ago and the invention of

movable type 600 years ago. In all these cases, the inventions involved *disintermediation*: the elimination of intermediaries in the transmission of knowledge. Writing eliminated the oral historians; one could read information from far away and long ago without having to speak to a person who had personally memorized that knowledge. Print allowed a far greater distribution of knowledge than handwritten books and scrolls, eliminating an entire class of scribes who controlled access to the precious and rare records. The 'Net and the Web have continued this trend, with a radical increase in the number of people capable of being publishers. Where publishing once required printing presses, capital, and extensive administrative infrastructure, or at least relatively expensive mimeographs (1950s), photocopiers (1960s), and printers (1970s), today an individual can publish material relatively inexpensively, if not free. Many Internet Service Providers (ISPs) offer free Web-hosting services and places for people to join electronic communities of every imaginable type. Even if the individual does not have access to the Internet at home, and follows work policies on use, free access can be obtained from local libraries or the increasing number of free Internet hot spots.

Web pages can lead to visibility unheard of even a decade ago. For example, one young exhibitionist named Jennifer Kaye Ringley put up a Website to display images of her home taken through Web-enabled cameras (Webcams); this "jennycam.org" site received up to *half a million hits per day* while it was in operation. Another young woman decided to put up a Website devoted to one of her favorite literary characters, Nero Wolfe, in the mid-1990s. Within a few years, her site was so well respected that she was hired by a Hollywood filmmaker as a technical consultant on a series of Nero Wolfe movies. The fees she was paid, despite offering to help for free, helped her get through her Ph.D. studies in social psychology. It would have been virtually impossible for her to achieve this recognition by trying to publish her own hard-copy fan magazine; the paper might have reached a few hundred people, but the Website reached many thousands.

Unfortunately, all of this disintermediation has negative implications as well as positive ones. Freedom from publishers has liberated the independent thinker from corporate influence, editorial limitations, and standards for house style. However, this freedom also liberated many people from responsible reporting, adequate research, and even the rudimentary principles of spelling and grammar. The dictum "Don't believe everything you read" is even more important when reading Web-based information. Individuals may publish incorrect versions of technical information (e.g., health sites that claim that massaging parts of the earlobe can cure many known diseases), unsubstantiated theories about historical and natural events (e.g., the Tungska Impact of 1908 was caused by an antimatter meteorite), and off-the-wall revisionist history (e.g., slavery in the United States was good for black people, and Hitler never persecuted Jews).

Wikipedia, although it has become the first line of information for some users, suffers from the possibility of temporary or even long-term modifications of content as pranks or for other purposes. For example, supporters of the failed right-wing candidate Sarah Palin are thought to have tried to alter the *Wikipedia* entry on Paul Revere to force it to conform to her incorrect statements about his famous ride.[5] Because there is no guarantee that the content of a reference to *Wikipedia* is legitimate or will be legitimate or even the same the next time the reference is followed, academic institutions tend to reject references to *Wikipedia* in term papers.

48.2.6 Hoaxes. Pranksters have been using email to fool gullible people for years using a particular sort of incorrect information: deliberate *hoaxes*. A hoax is a mischievous trick based on a made-up story. There are two major kinds of hoaxes circulating on the Internet: urban myths and false information about viruses. The

archives in the urban myths Websites are full of hilarious hoaxes, some of which have been circulating for years. Why don't they die out?

The problem is the nature of the Internet. Information is not distributed solely from a centrally controlled site; on the contrary, anyone can broadcast, or rebroadcast, any kind of data at any time. There is neither reliable creation dates nor obligatory expiry dates on files, so those receiving a five-year-old document may have no obvious way of recognizing its age, and they almost certainly have no simple way of identifying obsolete or incorrect information. All they see is that the document has been sent to them recently, often by someone they know personally.

48.2.6.1 Urban Myths. Here are some notorious examples of the bizarre and sometimes disturbing urban myths that are thoroughly debunked on the Snopes.com Website:

- **Expensive cookies.** Someone claims that a Neiman-Marcus employee charged $250 to a credit card for the recipe to some good chocolate chip cookies. This story has been traced to a false claim dating back to 1948 in which a store was accused of charging $25 for the recipe to a fudge cake.
- **Do not flash your car lights.** In a gang-initiation ritual, hoodlums drive down a highway with their car lights off. Innocent drivers, flashing their lights as a reminder, would become the new target victims, usually resulting in their deaths by the gang.
- **Watch out for poisoned needles.** Insane, vengeful druggies leave needles tipped with HIV+ blood in movie theater seats, gas pump handles, and telephone change-return slots.
- **Lose your kidneys.** The victim visits a foreign city, goes drinking with strangers, and wakes up in the morning in a bathtub of ice with two neat incisions through which both kidneys have been removed. No one ever seems to explain why criminals who remove kidneys would bother packing the victim in ice.
- **Poor little guy wants postcards.** Craig Shergold is just one of the many real or imaginary children about whom well-meaning people circulate chain letters asking for postcards, business cards, prayers, and even money. Shergold was born in 1980; when he was nine, he was diagnosed with brain cancer, and friends started a project to cheer him up—they circulated messages asking people to send him postcards so he could be listed in the *Guinness Book of World Records.* By 1991, he had received 30 million cards and an American philanthropist arranged for brain surgery, which worked: Shergold went into remission. The postcard deluge did not. By 1997, the local post office had received over *250 million* postcards for him, and he was long since sick of the whole project.
- **Wish you would stop Making a Wish.** Around the mid-1990s, some prankster inserted false information about the Make-a-Wish Foundation into the outdated chain letters concerning Shergold. The unfortunate organization was promptly inundated with email and postal mail, none of which was in any way useful or relevant to its work. They had to post disclaimers on the Website to try to dissociate themselves from the outdated information.

48.2.6.2 Virus Myths. One category of hoaxes has become a perennial nuisance on the 'Net: virus myths. There is something wonderful about the willingness

of gullible, well-meaning people to pass on ridiculous news about nonexistent viruses with impossible effects. One of the most famous is the Good Times "virus," which appeared around 1994. The myth and numerous variants have been circulating uninterruptedly for years. Every few years, there is a new outburst as some newcomers to the Internet encounter an old copy of the warnings and send it to everyone they know.

The original very short warning was as follows, including the incorrect punctuation:

> Here is some important information. Beware of a file called Goodtimes.
>
> Happy Chanukah everyone, and be careful out there. There is a virus on America Online being sent by E-Mail. If you get anything called "Good Times", DON'T read it or download it. It is a virus that will erase your hard drive. Forward this to all your friends. It may help them a lot.

The Good Times virus claimed that downloading a document or reading a document could cause harm; at that time, such a claim was impossible. Ironically, within a couple of years, it did in fact become possible to cause harm via documents because of the macro-language capabilities of Microsoft Word and other programs enabled for scripting. Over the rest of the 1990s, foolish people modified the name of the imaginary virus and added more details, sometimes claiming impossible effects such as destruction of computer hardware.

By 1997, the warnings were so ridiculous that an anonymous author distributed the following Monty Pythonesque satire:

> It turns out that this so-called hoax virus is very dangerous after all. Goodtimes will re-write your hard drive. Not only that, it will scramble any disks that are even close to your computer. It will recalibrate your refrigerator's coolness setting so all your ice cream goes melty. It will demagnetize the strips on all your credit cards, screw up the tracking on your television and use subspace field harmonics to scratch any CDs you try to play.
>
> It will give your ex-girlfriend your new phone number. It will mix Kool-aid into your fish tank. It will drink all your beer and leave dirty socks on the coffee table when company comes over. It will put a dead kitten in the back pocket of your good suit pants and hide your car keys when you are late for work.
>
> Goodtimes will make you fall in love with a penguin. It will give you nightmares about circus midgets. It will pour sugar in your gas tank and shave off both your eyebrows while dating your girlfriend behind your back and billing the dinner and hotel room to your Discover card.
>
> It will seduce your grandmother. It does not matter if she is dead, such is the power of Goodtimes, it reaches out beyond the grave to sully those things we hold most dear.
>
> It moves your car randomly around parking lots so you can't find it. It will kick your dog. It will leave libidinous messages on your boss's voice mail in your voice! It is insidious and subtle. It is dangerous and terrifying to behold. It is also a rather interesting shade of mauve.
>
> Goodtimes will give you Dutch Elm disease. It will leave the toilet seat up. It will make a batch of Methamphetamine in your bathtub and then leave bacon cooking on the stove while it goes out to chase gradeschoolers with your new snowblower.

Unaware people circulate virus hoaxes because they receive the hoax from someone they know. Unfortunately, a personal friendship with a sender is no guarantee of the accuracy of their message. Some awfully nice people are gullible, well-meaning dupes of social engineers. Transmitting technical information about viruses (or any apprehended danger) without verifying that information's legitimacy and accuracy is a disservice to everyone. It makes it harder for experts to reach the public with warnings of real dangers, and it clutters up recipients' email in-baskets with alarming

information of limited or no use whatever. Teach employees, family, and friends to consult snopes.com or an equivalent before forwarding alarming messages—and don't forward any messages about malware: Make sure everyone you know uses up-to-date antimalware software.

48.2.6.3 Junk Email. Unsolicited commercial email (UCE) is derisively known as *junk* email and also as *spam*. Junk email is spawned by foolish (in the early days) or criminal (today) people who send out thousands or millions of identical messages to unwilling recipients. Junk email clogs victims' in-baskets and wastes their time as they open these unwanted messages and take a few seconds to realize that they are junk. Junk email containing pornographic images or advertising pornography may be highly offensive to the recipients. Junk may even push people's email systems over their server limits if they are not picking up their messages regularly; in such cases, wanted email may bounce because the mailbox is full. Today, junk email is the primary vector for social engineering attacks such as *phishing* designed to trick recipients into compromising their privacy or their identification and authentication codes. (See Chapter 20 in this *Handbook* for more details of spam, phishing, and other tricks.)

Most junk email uses forged headers; that is, the senders deliberately put misleading information in the FROM and REPLY fields to avoid receiving angry responses from the victims of their criminal behavior. Forging email headers is illegal in the states of Massachusetts, Virginia, and Washington. In these states, if the perpetrators are identified, it can lead to court cases and financial penalties for each message involved in the fraud.

In one famous, groundbreaking case, college student Craig Nowak sent out a few thousand junk email messages and followed the instructions in his spam kit by putting a made-up REPLY address using "@flowers.com" without checking to see if there really was such a domain. Indeed there was, and the owner of this reputable floral delivery service, Tracy LaQuey Parker, was none too pleased when her system was flooded with over 5,000 bounce messages and angry letters from customers saying that they would never do business with her again. She sued the student for damages and was awarded over $18,000 by a judge who said he wished he could have been even more punitive.[6]

In general, spam has become a mechanism for tricking unaware vendors into paying for illusory marketing services. It is sad to see email advertising for Chinese industrial piping being sent to North American university professors; the victims are the hard-working Chinese industrialists who have been cheated by assurances from criminals promising to send their advertising to willing and well-qualified recipients.

Much of the remaining junk is sent out in the hope that a tiny proportion of the recipients of the misspelled, absurd claims will be gulled into sending money—or their email addresses—to drop boxes. For more information about such social engineering attacks, see Chapters 19 and 20 in this *Handbook*.

If you are involved in an email discussion group, especially an unmoderated group, about a specific topic, do not post email to members of the list on a subject that is outside the topic area. A typical class of inappropriate posting is an appeal for support of a worthy cause that has no, or only a tenuous, relation to the subject area. For example, someone might appeal for support to save whales in a discussion group about gardening: bad idea. The reasoning is "They like plants; probably environmentally sensitive; likely to be interested in conservation; therefore they will be glad to hear about whales." The problem is that such reasoning could be extended to practically any topic, disrupting the focus of the group. Such messages often cause angry retorts, which are typically sent by naïve members to the entire list instead of only to the sender

of the inappropriate mail. Then the angry retorts cause further angry responses about burdening the list with useless messages and soon the gardening group is mired in dissension and wasted effort, generating bad feeling and distrust.

As suggested in the preceding paragraph, if you see inappropriate messages on an email list you care about, do not reply to the entire list; reply *nicely* and only to the sender, with possibly a copy to the moderator, if there is one. The reply should be temperate and polite.

48.2.6.4 Chain Letters and Ponzi Schemes. A particularly annoying form of junk email is the chain letter. Some chain letters include ridiculous stories about terrible diseases and accidents that have befallen people who refused to forward the message (ridiculous on the face of it, but apparently appealing to the irrational enough to keep the nonsense circulating). Others focus on getting victims to send money to someone at the top of a list of names, while adding their names to the bottom of the list, before sending it on to a specified number of recipients. Depending on the length of the list and the amount to be sent to the person on top, the theoretical return could be in the hundreds of thousands of dollars. In practice, only the originators of the scheme profit. After a while, all possible participants have been solicited with disappointing results, and the chains are broken in many places.

Another type of pyramid is known as a Ponzi scheme, which is an investment swindle in which high profits are promised and early investors are paid off with funds raised from later ones. The scam is named after Charles Ponzi (1882–1949), a speculator who organized such a scheme in 1919 and 1920. The Ponzi scheme tricked thousands of people in Boston when Ponzi guaranteed a 50 percent profit on contributions in 45 days and a doubling of value in 90 days. The con man claimed he was redeeming 1-cent Spanish postal certificates for 6-cent U.S. stamps—a claim ridiculed by financial analysts at the time. Nonetheless, Ponzi took in around $15 million in 1920 dollars and stole around $8 million, paying out the rest to the early participants in order to develop credibility. Six banks collapsed because they invested their depositors' funds in the scheme. Ponzi eventually served over three years in jail but escaped in 1925.[7]

The modern-day email Ponzi scheme typically includes passionate assurances from vaguely identified people about how skeptical they were about the scheme, but how they succumbed to curiosity, participated in the scheme, and earned vast amounts of money (e.g., $50,000) within a couple of weeks. The letters often include assurances that everything is legal and point to nonexistent postal information phone lines or claim "As Seen on TV" at various points in the letter.

These letters instruct the victim to send a small amount of money (typically $1 or $2) to a short list of about four people to receive their "reports." The victim is then instructed to add his or her name and address to the list, while removing the first one, before sending a copy of the new letter to as many people as possible. Some letters go through computations involving such assumptions as "Imagine you send out a hundred, a thousand, or ten thousand messages and get a mere 1%, 2%, or 10% response," and then promise enormous returns. In fact, the "reports" are nothing but one-page, meaningless blurbs about chain letters. The scammers are trying to get around regulations such as the U.S. Post Office's bar against fraudulent uses of the mail.

Here is the exact text of a letter sent on December 1, 2000, by V. J. Bellinger of the Operations Support Group of the United States Postal Inspection Service in Newark, New Jersey. It has some interesting information that should be helpful to

readers attempting to convince employees (or family and friends) that such chain email involving postal addresses is illegal.

> A chain letter or a multi-level marketing program is actionable under the Postal Lottery, False Representation, and/or Mail Fraud Statutes if it contains three elements: prize, consideration and chance. Prize is usually in the form of money, commissions, or something else of value that the solicitation claims you will receive. Consideration is the required payment to the sponsor in order to obtain the prize. Chance is determined by the activities of participants over whom the mailer has no control. These types of schemes constitute lotteries are barred from the mails because they violate the following statutes: Title 18, United States Code, Sections 1302 and 1341 and Title 39, United States Code, Section 3005.
>
> In attempts to appear legal, many chain letter or multi-level marketing mailings offer, for a fee, a product, or "report." However, since the success of the program is dependent on the number of people willing to participate, all three elements that constitute a violation continue to be present.
>
> The promoter of this scheme has been advised of the potential violations involved and has been requested to discontinue this type of mailing activity. …

A superficially similar phenomenon is known as multilevel marketing. In this nonfraudulent, legitimate system of selling products and services, people are encouraged to recruit distributors from among their friends and acquaintances, but the emphasis is on the value of the products. No one claims that anyone is going to become wealthy without work, and there is no demand for investments. The products have an established market, and the company makes money through sales, not through recruitment.

Here are some practical guidelines for employees and individuals:

- Do not participate in any scheme that relies on forwarding large numbers of letters or email messages to everyone you know or to strangers.
- Differentiate between pyramid frauds and legitimate multilevel marketing systems: The former emphasize enrolling participants, whereas the latter emphasize the value of products and services.
- Do not participate in alleged multilevel marketing systems if they require substantial investments.
- If you are interested in a multilevel marketing operation:
  - Check out the owners and officers.
  - Talk to people who have bought the products to see if they are happy with their purchases.
  - Contact your local Better Business Bureau to see if there have been any complaints.
- Do not send money to suspected pyramid frauds.
- Work with your colleagues to demonstrate how a pyramid fraud takes money from a growing number of later victims and shifts it to people who participate earlier in the fraud. Reinforce the fact that fraud is illegal, even though the prospect of early participation might seem to yield results.

48.2.6.5 Get-Rich-Quick Schemes. Other get-rich-quick schemes on the 'Net play on the victims' wishful thinking, their lack of skepticism, and usually on a lack of common sense. There have been claims that you can earn a quarter of a million dollars a year by grooming poodles in your home. Or that you can become a millionaire

by working four hours a week, sending out promotional literature for products you do not even have to sell. Often, dangerous people promulgate some such schemes; for example, some extremist militia groups have been charging people hundreds of dollars to learn how to defraud the government by claiming liens on government property and then pledging the nonsensical liens as collateral for loans. Other criminals circulate programs for generating fraudulent credit card numbers and using them to steal goods. In other cases, criminals charge money to teach victims how to falsify their bad credit records so they can obtain yet more fraudulent credit, all the while claiming that their criminal methods are 100 percent legal.

From a corporate standpoint, such chain letters and schemes waste bandwidth and pose a potential for serious embarrassment when enterprise resources are used to spread the nonsensical material. However, corporate security can win favor with users by helping them avoid the pitfalls of such fraud, even when using their own computer systems. The benefits are particularly strong when helping employees to teach their own children how to avoid this kind of trouble.

To illustrate the trouble kids can get into using these techniques, consider the case of Drew Henry Madden. In 1996, this 16-year-old Australian boy from Brisbane, just after leaving school, started defrauding businesses using stolen and forged credit card numbers. He stole $18,000 of goods and, in February 1997, pled guilty to 104 counts of fraud and was sentenced to a year in jail. However, investigators uncovered additional fraud, and it turned out that he had stolen an additional $100,000 in goods and services. In October 1997, he pled guilty to another 294 counts of fraud and was given an additional suspended sentence. His defense attorney blamed poor security for the losses: "Madden started with very minor credit card fraud, but it escalated alarmingly, because the safeguards were so inadequate." Despite the youngster's unusual revenue stream, his mother appeared to have accepted his globetrotting ways and massive purchases of lottery tickets without comment. At one point, she told reporters, "If we were a wealthy family he'd be at a private school, where his talents could be directed properly."

A relatively new kind of fraud on the Internet is the diploma mill. These organizations pretend to be educational institutions; actually, they are one or more fraudulent individuals who sell bogus diplomas purporting to represent recognized degrees but that fool no one but the purchaser. While diploma mills are not accredited, the lack of accreditation does not automatically implicate a school as a fraudulent entity.

48.3 THREATS TO PEOPLE AND SYSTEMS. One particular class of email deserves a special mention: threats. Threatening emails may target people, systems, organizations, or the processes that these entities rely on. As with the other types of illegal activity on the Internet, proper education, awareness, and response can limit the number of future victims.

48.3.1 Threats of Physical Harm. Anyone who receives threats through email has a right, and possibly a duty, to inform local law enforcements officials. In today's climate of fear and violence, any threat warrants attention. In addition to the distress such messages can generate, they may be warning signs of serious trouble. In particular, threats about violence at work, at school, or against any definable group may be the early warning that allows authorities to step in to defuse a potentially explosive situation.

Sending threatening email is not an acceptable joke or a minor prank, especially if the threat involves violence. Some people, believing that they can mask their real identity,

have foolishly sent death threats to the White House; because the Secret Service is obligated by law to investigate *all* threats to the president and the first family, agents show up within a few hours to interrogate the miscreants. For example, youngsters in the tenth grade at Profile High School in Bethlehem, New Hampshire, sent death threats to the White House Website from their school computers. The messages were traced within minutes by the Secret Service, the children were suspended from school, and they lost their Internet privileges for the next two years.

48.3.2 Pedophiles Online. This section applies primarily to training users for protection of their children.

Pedophilia is defined as sexual arousal in response to contact with or images of prepubescent children. Some pedophiles misrepresent themselves as youngsters in chat rooms or via email and trick children into forming friendships with what they believe are peers. In one notorious case, Paul Brown Jr., a 47-year-old man, misrepresented himself as a 15-year-old boy in email to a 12-year-old girl in New Jersey. The victim's mother stumbled onto the long-range relationship when she found sitting on her own doorstep a package from her daughter to a man she did not know; the child had put the wrong postage on it and the post office had sent it back. Opening the package, she found a videotape that showed her daughter cavorting naked in front of the family video camera. The distraught mother searched her daughter's room and discovered a pair of size 44 men's underpants in one of the child's bureau drawers.

Brown was arrested in February 1997. Police found correspondence with at least 10 other teenage girls across the country, through which Brown convinced his young victims, some as young as 12, to perform various sexual acts in front of cameras and to send him the pictures and videotapes. He pleaded guilty in June to enticing a minor into making pornography. In August 1997, at his sentencing hearing, one of his many victims told the court that she had suffered ridicule and humiliation as a result of her entrapment and had left her school to escape the trauma. She accused Brown of emotional rape. Displaying an astonishing interpretation of his own behavior, Brown said at his sentencing hearing, "It was just bad judgment on my part." Using good judgment, the court sentenced him to five years of incarceration.

In March 2000, Patrick Naughton, a former executive of the INFOSEEK online company, pled guilty to having crossed state lines to commit statutory rape of a child. In August, FBI officials said that Naughton had been providing help in law enforcement investigations of pedophilia on the 'Net. In return for his cooperation, prosecutors asked the court for five years of probation (instead of a possible 15 years in prison), counseling, a $20,000 fine (instead of the maximum $250,000), and an agreement not to have unapproved contact with children and to stay out of sex chat rooms online.

The problem of Internet-enabled pedophile stalking has reached international dimensions. In January 1999, police forces around the world cooperated to track and close down a worldwide ring of pedophiles trafficking in child pornography through the 'Net. Child safety experts have warned the U.S. congressional committee on child online protection that with the average age of online users declining (children between the ages of two and seven are among the fastest-growing user cohorts on the Internet), children increasingly are put at risk by their careless or ignorant online activities.

48.3.3 Viruses and Other Malicious Code. As of 2008, the WildList (www.wildlist.org) reports over 2,000 distinct forms of malicious program code commonly circulating in cyberspace. There are many more types recorded by antivirus researchers, but they have not been seen infecting significant numbers of user

computers. Most of these harmful programs are limited to antivirus laboratories and to the computers of virus hobbyists—people who derive pleasure from playing with danger. For more information about viruses and other malware, see Chapters 16, 17, 18, and 41 in this *Handbook*.

Employers should have policies clearly forbidding the creation, exchange, and storage of malicious software on corporate systems.

48.3.4 Spyware and Adware. In December 1999, computer scientist, cybercrime investigator, and writer Richard Smith became curious about a program called *zBubbles* that he had installed on his system to improve online shopping. Created by Alexa, a subsidiary of e-tailer Amazon.com, the program provided competitive information about alternative and possibly cheaper sources for particular products. However, Smith discovered that there was more going on than met the eye.

Smith monitored his own Internet traffic while he was using zBubbles by using a packet sniffer, a tool that displays details of every piece of information being transmitted through a network connection. He found that zBubbles was sending a steady stream of information about him and his surfing habits to Alexa, including his home address, the titles of DVDs he had browsed on Buy.com, and the details of an airline ticket he had verified online. In addition, the program even continued to send information regularly to Alexa's servers even when Smith was not using his browser. It was learned that zBubbles was not the only program sending information back to its makers.

Many programs are available that, once installed, report on the Websites you visit, which banner advertisements you click, what products you search for, and any other information the programs have been designed to acquire. Even widely used downloading software, such as NetZip, has been shown to report to its providers on the names of every file downloaded by each user.

Sometimes these programs are informally known as *E.T.* applications, in a reference to Steven Spielberg's movie of that name, in which an extraterrestrial strives to "phone home"—exactly what the spyware programs are doing.

The term *spyware* is applied to any technology that transmits information without the knowledge of its user. Several programs distributed without charge through the Internet secretly collect information about the user, monitor user behavior, and then send those data to advertisers. The more general class of monitoring software that collects information for use by advertisers is known as *advertising-supported software* or *adware.* These programs allow freeware to make money for its creators by generating revenue based on how many users transmit information to the advertisers about their habits.

Although defenders of the advertising-supported programs claim that they are harmless, privacy advocates argue that the issue is control: Do users know what these programs are doing, or are they collecting and transmitting information covertly? Some adware comes with complicated contracts containing complex legal language to bury the fact that they will monitor and report user behavior. Worse yet, many such contracts explicitly authorize the software supplier to alter the privacy conditions without notification and, preposterously, instruct the user to check the contracts on the Web frequently. No one has the time to monitor countless suppliers to see if privacy conditions have been altered, especially if there is no attempt to highlight changes.

Another issue is that some spyware modules have used *stealth* technology characteristic of viruses, Trojan horses, and other malicious software. For example, some adware (e.g., TSADBOT) installs itself as a system process and is not listed in the Windows task list. Therefore, it cannot easily by aborted by a user. TSADBOT also

resists removal; even if the carrier product is uninstalled, TSADBOT persists. If a user's firewall blocks outbound transmission by the TSADBOT process, the spyware initiates attempts to reach its target at a rate of 10 per second, potentially leading to central processing unit (CPU) and network resource overload.

Spyware, like any software, can contain errors that cause system problems. In particular, components of the Aureate/Radiate spyware have been shown to cause system instability and crashes.

One of the most egregious cases of spyware erupted in 1999, when it was discovered that CometCursor, a supplier of cute cartoon-character cursors aimed at children, was sending information back to its servers about what the children were browsing on the 'Net. According to some attorneys, this kind of covert data gathering about children may be a violation of the U.S. Federal Child Online Privacy Protection Act.

Several free software programs have been written to help users identify and remove spyware. In addition, personal firewalls can usually identify and block unauthorized outbound communications; the free version of ZoneAlarm, for example, does so effectively. Today's antimalware programs (e.g., Bitdefender) include antispyware functions. There are also many specialized programs (e.g., Lavasoft's Ad-Aware) available that run in the background to monitor and thwart attempts to install spyware and adware.

48.4 THREATS TO PRODUCTIVITY. Some activities and phenomena are nuisances to employers principally because of their noxious effects on productivity and their abuse of corporate resources. Junk email and mailstorms, for example, are a problem because they saturate resources, not because they cause specific harm to the organization or to its employees. However, chain letters, get-rich-quick schemes, online auctions, online gambling, excessive online shopping, and Internet addiction can be directly harmful to employees and others.

48.4.1 Inefficient Use of Corporate Email. The next sections focus on problems caused by mistakes in the use of email—mistakes that can cause annoyance, inefficiency, and potential disruption of critical business processes.

48.4.1.1 Forwarding Email to Personal Accounts. Employees may be tempted to forward their corporate email traffic to their personal email addresses for convenience, or when they have no convenient way of accessing their corporate email system from outside the office. Such forwarding should be forbidden by policy unless virtual private networks (VPNs) or other strongly encrypted channels are used for the employee's private email.

Email and other traffic on the Internet have no inherent confidentiality. In theory, anyone capable of intercepting TCP/IP packets anywhere during transmission can breach confidentiality. Thus, again in theory, anyone with access to the equipment of ISPs, Internet backbone transmission lines, and even to the public switched telephone network can intercept packets. With downlink footprints from satellite relays amounting to square miles, practically anything can in theory be intercepted from much of the traffic circulating on the Internet.

However, in practice, reported breaches of confidentiality have almost all resulted from data access at the endpoints, not in transit. Insider attacks and breaches of server security have been responsible for most of the data interceptions that have reached the press and the courts.

A practical impediment to effective interception of meaningful data in transit is the datagram routing that underlies the Internet: Datagrams are packets of information

with origin and destination information; store-and-forward transmission allows these datagrams to be sent through the Internet via different routes from other packets in a message stream. Routing tables can be updated in real time to reflect changes in traffic density or availability of specific links to other destinations on the Internet, so there is no guarantee that packets from the same message will travel the same route or arrive in the proper sequence (sequence numbers allow reassembly of the original message). Therefore, seizing individual packets at random anywhere other than the origin and destination of packets is unlikely to result in very much result for the effort.

Nonetheless, best practices do recommend that encryption be used for communication of sensitive data; therefore, many organizations install virtual private networks (VPN) for communication with established trading partners. VPN software is also available for *tunneling* through the Internet from a remote workstation over nonsecure communications lines. A simple example of such a link-encryption function is the Web-based email services that use SSL to establish a secure link to the email server (i.e., they use https instead of just plain http). The user can pick up email from the corporate server without having it forwarded in the clear to an insecure external email service. Some of the email products include facilities for direct communication between a secure email server and the users' email client.

Using *VPN tunneling software* as a search string in the Google search engine brings up almost half a million hits (in May 2013), many of them for specific products and data sheets, so readers will be able to find a solution that fits their needs.

48.4.1.2 Mislabeling the Subject Line. Many people make the mistake of creating new messages to a correspondent by finding any old message from that person and replying to it. The problem is that these people usually leave the old subject intact, resulting in ridiculous situations such as finding a critically important message in July in an email labeled "Birthday party 12 May."

Not all email messages are created equal; some are destined for the trash heap, if not of history, at least of the email system. That decision is sometimes made automatically as a function of the subject line. For example, a user adds the subject line of a joke to an email filter, resulting in future messages with that subject ending up in the junk mail folder. Someone replies to the joke message with important information, and the mail filter sees the subject and automatically moves the message to the recipient's junk mail folder. The recipient may never see the important information, as most people do not actively monitor their junk folders.

Another problem with mislabeled subjects occurs when someone embeds more than one distinct topic in an email message whose subject line implies otherwise. For example, suppose an email message subject reads "Next week's meeting," but the sender includes an urgent request for action today on some critical issue; there is a good chance the receiver may not open the message right away if other messages seem more important.

Employees should make their subject line as descriptive as possible without turning it into a paragraph. Some email systems truncate subject lines in the display of messages that a user sees; it makes sense to put keywords at the front of the subject. Encourage staff to use prefixes such as "MISA:" or "ABCv2.0.1:" to help organize their messages. Using standard formats in subject lines can help too. For example, faculty and staff in the MISA program at Norwich University refer to an issue in a particular seminar by using the form "MISA c.s" in their subject line, where c represents the class (e.g., 40 for students starting in December 2013) and s represents the seminar number (e.g., 1 through 6).

These simple suggestions can make email more effective as a communications tool.

48.4.1.3 First e-Impressions. When you receive an email message from a stranger, do you care whether it has spelling mistakes and grammar mistakes? What about offensive language and off-color humor? Does the context matter? For example, do you apply the same standards to email referring to business matters as to informal communications about a hobby?

Researchers at the University of Chicago have been investigating the effects of email on perceptions of character. Psychologist Nicholas Epley and colleagues examined oral exchanges on conversational topics by phone between randomly selected people using six assigned questions. They then transcribed the oral conversations and used exactly the same answers for the written, email version of the question and answer sessions.[8]

Their results were interesting. The questioners had been given false biographical sketches of the people they were communicating with, indicating substandard intelligence or normal intelligence, as well as different pictures showing neat people or slobs. Subjects who used the phone to listen to the prescribed responses had favorable impressions of their interlocutor's intelligence, regardless of the bios and pictures. In contrast, "Via email, however, students held onto their first impressions, continuing to assume their partners had substandard intelligence, for example, if that's what the biographical sketch indicated."

If this research is confirmed, the lesson is that when using email, first impressions really do count. Professionals should carefully review email messages for acceptable writing, including word choice, punctuation, capitalization, and spelling.

48.4.1.4 Email Disclaimers. Author Kabay once received a 30-word email message from a very nice reader in Britain and noticed that his email system added the following astonishing disclaimer, which is quoted in its sonorous totality, including British spelling, after scrubbing it of identifying details:

> This email, its contents and any files or attachments transmitted with it are intended solely for the addressee(s) and may be legally privileged and/or confidential. Access by any other party is unauthorised without the express written permission of the sender.
>
> If you have received this email in error you may not copy or use the contents, files, attachments, or information in any way nor disclose the same to any other person. Please destroy it and contact the sender on the number printed above, via the <Name of Bank> switchboard on +44 (0) nnnn nnnnnn for <place1> and + 44 (0) nnnn nnnnnn for <place2> or via email by return.
>
> Internet communications are not secure unless protected using strong cryptography. This email has been prepared using information believed by the author to be reliable and accurate, but <Name of Bank> makes no warranty or representation, express or implied, as to its accuracy or completeness and is not liable to you or to anyone else for any loss or damage in connection with any transmission sent by the Bank to you over the Internet. <Name of Bank> makes no warranty that any information or material is free from any defects or viruses.
>
> In particular <Name of Bank> does not accept responsibility for changes made to this email after it was sent. If you suspect that this email may have been amended or intercepted, please contact the sender in the manner stated above. If this transmission includes files or attachments, please ensure that they are opened within the relevant application to ensure full receipt. If you experience difficulties, please refer back to the sender in the manner stated above.
>
> Any opinions expressed in this transmission are those of the author and do not necessarily reflect the opinions of the Bank and may be subject to change without notice.
>
> Please note that for the purposes of this document all references to <Name of Bank> or the Bank shall be taken to mean <Name of Bank> (place) Limited or any other member of the <Bigger> Bank Group. Nothing in this transmission shall or shall be deemed to constitute an offer or acceptance of an offer or otherwise have the effect of forming a contract by electronic communication.

Kabay commented in his response, "Did you know that your message has 30 words (152 bytes including spaces) whereas your disclaimer has 367 words (2,177 bytes)? That's the lowest signal-to-noise ratio (6.5 percent useful information out of the total and a 1:73 signal:noise ratio) I've ever seen outside a copy-of-copy-of-copy chain. Please congratulate your attorneys on using maximum of bandwidth for minimum content!"

Cluttering up email messages this way is a waste of bandwidth. It is worse in offices where people copy entire messages without editing the contents, resulting in copy-of-copy-of-copy chains that spread like cancerous eruptions through in-baskets throughout the organization. Some well-meaning folks even include the detailed headers in their copies.

As a matter of courtesy and good sense, when one replies to a message, it is a simple matter to strip nonessentials out of the copy of the original. Senders can use ellipses (… for cuts within a sentence, … . for cuts crossing sentence boundaries) to signal gaps, but usually one or two snips are enough to clean up the copy so that the reader can get the gist of the conversation without having to wade through reams of superfluous stuff. Unfortunately, this recommendation does not seem to be used much in practice.

48.4.1.5 Centralized Distribution Lists. Organizations may grow large enough that there is significant turnover among the staff. Not only do new staff members periodically join the group, but also staff members move from one functional group to another; for example, a staff member may change from being an assistant director in one program to being an administrative director in another. Occasionally, staff members may leave the group altogether.

A primitive way of maintaining distribution lists is to name a "Keeper-of-the-Lists" to maintain the list of all staff members; however, there is no link between the file and the mailing lists that each member of the group must maintain to be able to distribute email to appropriate individuals or groups. The independent files are almost certain to diverge from a centralized and accurate list. For example, a message that should be sent to all current employees may end up missing several new members and including staff members who no longer work in the target group.

Trying to make many people maintain their own copies of several distribution lists is a hopeless cause: Even with the best will in the world, people will inevitably forget to update their lists and therefore:

- Some mailings will miss legitimate recipients.
- Some people will receive messages they have no business reading.

There are at least four solutions that would rectify such a problem.

1. One can implement a central email sever (e.g., Microsoft Exchange Server), switch all users to a centrally controlled email client (e.g., Microsoft Outlook), and define corporate distribution lists maintained by the Keeper-of-the-Lists. All users will automatically access the one and only distribution list for each group without manual intervention.
2. One can install widely available list-server software to allow centralized creation and maintenance of specific lists; for example, SGS-ALL, SGS-DIRECTORS, MSIA-STAFF, MSIA-INSTRUCTORS, and the like create lists that all employees can use in addressing email.

3. One can switch all users to any email client that supports exportable mailing lists. Updated corporate distribution lists can then be sent to all users. However, this solution still requires manual intervention by users: Everyone has to replace an old list by the new list.
4. One can create a discussion group on a public server (e.g., Yahoo Groups) to define closed groups. These groups provide automatic access to mailing lists. Unfortunately, this approach has serious problems:
   - There are security concerns about using such groups for corporate communications.
   - It seems inappropriate to put a necessary production application on a free resource completely out of the control of the organization.

48.4.1.6 HTML Email. One of the six fundamental attributes of information that we protect is integrity, one aspect of which is consistency with the originally stored data (see Chapter 3 in this *Handbook*). When someone goes to the trouble of producing an elegantly formatted memorandum or other document and sends it out to recipients, everyone would like to preserve data integrity by seeing the same appearance on all the systems sharing that document.

Unfortunately, sending formatted messages as email messages (as distinct from attachments) does not guarantee preservation of the exact appearance of the source material.

Attractive, well-formatted email messages with boldface, italics, different point sizes, and the like usually get transmitted as HTML (hypertext markup language) to recipients' mailboxes, where most people's email clients (Eudora, Netscape, Outlook, etc.) allow the funny-looking code to be reconstituted into something similar to the original.

The word *similar* is mentioned rather than *exactly like* because HTML does not necessarily control the final appearance of text on a recipient's system. The codes refer to types, not exact matches, of fonts; thus, a sender might want to use, say, 24-point Arial as a Heading 1 display but a particular recipient might have defined Heading 1 as, say, Times Roman 14 point. A two-page original document may appear to be a three-page document to one recipient and a one-page document to another.

More significantly, though, many people turn off HTML email for security reasons. All such formatted email gets converted automatically into plain ASCII text. A correspondent once sent author Kabay a message that read: "Note: The on-line course evaluation system may be used from room, lab, and home—anywhere Internet access is available. / Overview: … . Failure to complete a course evaluation will result in a 'hold' being placed on the student's final grades."

The fragment of message that follows shows the result of MS-Outlook autoconversion of the original formatted HTML message to ASCII: "Note: The on-line course evaluation system may be used from room, lab, and home ? anywhere Internet access is available./Overview: … . Failure to complete a course evaluation will result in a ?hold? being placed on the student?s final grades."

- In the conversion process, the original apostrophes turned into question marks ("?hold?") because the sender was using "curly" quotation marks instead of the straight ones in the word-processing package or email editor. If one cares to prevent this peculiarity when using earlier versions of Microsoft Word, one has to turn off the option in the {Tools | AutoCorrect | AutoFormat As You Type}

screen by unchecking the box labeled {"Straight quotes" with "smart quotes"}. In later versions, turn off the option by clicking {Word Options} in the main menu, selecting {Proofing}, and then clicking the {AutoCorrect Options} button at the top of the box. {AutoFormat As You Type} is one of the available tabs, and you can then uncheck the box labeled {"Straight quotes" with "smart quotes"}.

- In addition, it looks like a dash character may have been in the text in the first section (labeled "Note"). One can turn that conversion off in the same menus by unchecking {Hyphens (–) with dash (—)}.

A much simpler solution to prevent the mess is simply to send unformatted ASCII text in all outbound messages by selecting that option in one's email package.

Some people try to send files that should look the same on a recipient system and the originating system by attaching word processing documents: for example, Word (DOC) files, WordPerfect (WPD) files, or Rich Text Format (RTF) files (and so on). Unfortunately, even these attempts do not necessarily work as planned, since lack of shared fonts, different default paper sizes (different countries may use different sizes), and different printing margins (resulting from installation of different printers) may cause the documents not to look precisely the same on all systems.

So if the exact appearance of a message one is sending via email is critically important, one should send the content *and* its format in a way that is (largely) platform independent; for example, Acrobat PDF (Portable Document Format) files. Although even they do not necessarily result in perfect rendition of the author's intentions across systems, PDF files are far more likely to succeed than the other methods mentioned. One can create PDF files in a number of ways; some systems have Adobe Acrobat installed so that one can either send to an Acrobat driver to create the PDF files or even just click a toolbar button to do so from within the word processor. Microsoft Office 2007 and later versions, for example, provide the ability to Save as PDF in all of its major components. Other packages exist that are less expensive (and generally less feature-rich) than the full Adobe Acrobat software, but nonetheless allow users to create PDF files easily. One can type "create PDF" into a Web search engine to find lots of choices.

48.4.1.7 Distribution Lists in Email. As for confidentiality, consider that using the *To* and *CC* (*carbon copy*—a bit of historical detritus) fields in email makes all recipient addresses visible to all recipients. This situation is usually helpful in internal email because team members can see who has gotten the message, but it can be annoying in external email. Why should a list of dozens, or even hundreds, of names of strangers be distributed freely among them, without the explicit permission of all concerned? Who knows where that information will end up? Use of the *BCC* (blind carbon copy) field eliminates the ability of recipients to see all of the intended recipients for the original message. This extra step is a good piece of email etiquette, whether the message is business or personal. The *BCC* field is also useful for internal email when the list of recipients is very large, but it is not important for people to know exactly who received the message.

In one case, a nice lady in the human resources (HR) department at a university sent out a note to a dozen people reminding recipients that they had not yet finished signing up for their new medical insurance coverage.

Unfortunately, she put all the email addresses into the *CC* (carbon copy) line where they were visible to everyone in the list. Predictably, someone on the list composed a

response to her, hit *REPLY ALL*, and sent some mildly personal information about the state of her medical concerns to all the recipients on the original list, none of whom had any interest in her problems.

Luckily, there was not a lot of private information in that message, but it did prompt the realization that many people unthinkingly use the *CC* line for addresses to a distribution list and that many people unthinkingly use REPLY ALL for replies to every email message.

The combination can lead to embarrassing violations of confidentiality; when the HR department staff use *CC* instead of *BCC* (the Blind Carbon Copy function that conceals the distribution list), the REPLY ALL function can inadvertently violate privacy.

In this case, there was no particularly sensitive material revealed, but a different case could easily violate HIPAA (Health Information Portability and Accountability Act) and the university's rules on employee confidentiality.

Once employees understand the issue, they will learn not to use *CC* for distribution lists when the intention is to communicate with individuals; by default, everyone should use the *BCC* list unless there is a need to stimulate group discussion of an issue or it is important for the members of the group to know who received the message.

It is important not to dismiss this issue as too easy or too obvious to bother with. "Against stupidity, the gods themselves contend in vain," wrote Friedrich von Schiller in his *Maid of Orleans* (*Die Jungfrau von Orleans*) in 1801. Nonetheless, the *CC* + REPLY ALL habit becomes a covert channel for release of confidential information for people who refuse to keep an address book, and simply look up any old email and REPLY ALL to it as a lazy way of sending a new message.

If you doubt the seriousness of the problem, take some time to look through your own archives of email and count how many obvious cases there are of emails with inappropriate subject lines and inappropriate distribution lists sitting in your received folders. Unfortunately, you may be dismayed by the results of your research. If you look into your own SENT folder, you may be even more dismayed.

48.4.1.8 Effective Use of BCC. As discussed, the problems caused by *CC* are worse when the recipients do not know each other. One often receives messages from technically unsophisticated correspondents who put dozens of email addresses in the *CC* field even though many of the recipients are total strangers to each other. Such exposure of email addresses always makes security staff nervous; who knows whether everyone on the list is trustworthy? Even if the list is not misused for outright spam, people often REPLY ALL with useless information, effectively adding people to discussion lists that they never wanted to be on.

One particularly annoying habit is to REPLY ALL with a comment stemming from some initial message. People then generate a series of increasingly long messages including copies of all the previous copies of the ostensibly clever repartee, driving some users to generate an addition to their junk mail filters.

The habit of using REPLY ALL is annoying enough when a reply does not in fact have to go to everyone on the original distribution list. However, REPLY ALL is a positive menace if it is coupled with the abhorrent practice of using an existing email message as a shortcut to creating a new one with a completely different topic, as discussed in Section 48.4.1.2.

48.4.1.9 Managing Private Email at Work. What is wrong with using corporate email for jokes, invitations, and the like? One issue is the waste of bandwidth. Some people find the quality of the jokes, hoaxes, and cheering sessions low enough

to be irritating. Worst yet, the tolerance level for what is considered appropriate in the workplace may vary by individual, requiring the utmost care and consideration for everyone. Another problem arises with politically sensitive messages, such as announcements or viewpoints that some members of a group may find offensive. Why should everyone in the group be subjected to a barrage of unsolicited email just because they work somewhere?

The question also raises some valuable and instructive points about appropriate-use policies for email. Corporations must have a formal written policy on appropriate use of official email. Managers should frame clear written policies that any of the staff members can easily consult for guidance about suitable and unsuitable content for personal messages using corporate mailing addresses. Such policies will reduce possible disappointments and resentments resulting from decisions based on unwritten expectations. In addition, any hint of discrimination based on particular political or religious biases will have to be scrutinized to ensure that the organization is not subject to legal repercussions.

An easy tool that employees can develop is a voluntary mailing list of nonwork email addresses for nonwork email. A Yahoo! group (http://groups.yahoo.com/), for example, offers many benefits over an informal list in the *CC*: or To: field. Jokes and the like can thus be distributed only to willing recipients, since joining can be purely optional. However, employees must always remember that any activities occurring on company equipment, or using company computing resources, may be viewed by authorized parties and could potentially hurt their reputation or, even worse, set them up for legal problems.

48.4.2 Mail Storms. A peculiar kind of junk email is sent by accident. These flurries of unwanted messages are called *mail storms.*

Most of us belong to mailing lists; many of us have more than one email address; some of us use autoforwarding to shift email from one address to another automatically; and a few of us use automated responses on our email accounts to let correspondents know when we are out of the office or unable to respond quickly.

All of these factors can contribute to mail storms.

48.4.2.1 Autoforwarding. A mail storm occurs when computers begin sending mail to each other without human intervention. Sometimes mail storms can become a denial of service by saturating communications channels and other resources. The email-enabled worms such as Melissa, the I-love-you message, and others, are examples of malicious software programs whose authors deliberately wrote them to create mail storms.

A simple situation occurred in the 1990s:

- An employee leaving on vacation decided to receive company email using a personal account on an ISP with a global presence. By setting an autoforward command on the company account, all incoming mail was sent to the personal email account.
- Unfortunately, on the remote tropical island where the vacationer spent two weeks, it was impossible to access the worldwide ISP without paying a surcharge of $6 a minute for long-distance service to the nearest dial-up node. This proved too expensive, and no emails were received or sent.

- Meanwhile, the company account dutifully forwarded every message it received to the proper personal account—which had a tiny storage limit of 250 messages. That limit was reached within a few days. At that point, every inbound message generated a bounce informing the sender that the recipient's mailbox was full.
- The very first full-mailbox message sent to the company account was autoforwarded back to the vacationer's personal mailbox.
- That copy of the full-mailbox message generated a second mailbox-full message, which then got bounced back to the company account, and so on without letup.
- Eventually, even the company mailbox filled up, and then the two email systems continued chattering at each other indefinitely. In this particular case, system administrators noticed the problem when the user's mailbox reached 20,000 messages and crashed the mail server.

The number of email messages that can be generated by this kind of infinite loop is a function of the latency of the positive feedback system that the user has accidentally created. For example, if it takes exactly one minute for a bounce message to be returned to the originating site, then each message causing an initial error can create 60 additional messages per hour. However, every new message from another sender that arrives at the originating mailbox will generate its own new set of bouncing messages in infinite loops. It is not uncommon to see tens of thousands of messages accumulating in the recipient's mailbox if nobody notices the loops with traffic mounting steadily into hundreds or thousands of messages per hour bouncing between the accounts, potentially generating a denial of service on corporate email through bandwidth saturation alone. The mail servers may also crash because of the overwhelming traffic.

An out-of-office message can also inadvertently create mail storms through a *race condition* (see Chapter 39 in this *Handbook*). For example, two employees (Albert and Bob) both enable out-of-office messages, and Albert sends an email to Bob. Bob's email service sends Albert back its out-of-office message, which in turn generates another out-of-office message from Albert to Bob. A mail storm results.

48.4.2.2 Poorly Configured List Servers. The user of an autoresponder may belong to a list where the FROM address is actually the broadcast address that sends a response to the entire list. The very first automated out-of-office response to the list will generate a message to everyone on that list, producing an infinite sequence of to-and-from messages. This situation is very embarrassing for the list administrator and intensely annoying for everyone else.

48.4.2.3 Human Error. Something analogous to a mail storm results from thoughtless behavior when using a public list. A typical instance occurs when a list member posts to an entire list comments relevant only to one individual. For example, a member asks for a reprint of an article and another answers on the list: "I'll send you a reprint tomorrow." Several thousand unwilling readers now know about this projected email message. One of these irritated people posts a message saying, "Did you really have to post that message to the entire list?" This second message is so irritating that at least one other person posts a third message to the entire list, criticizing the originator of the second letter for criticizing the writer of the first. This useless tempest of email continues via the public list, creating thousands of copies of useless information.

Another form of inconsiderate behavior is to quote entire messages when responding to email. Only the fragments of text that have elicited a response should be copied to the

new message. This principle is particularly important on public lists, where messages have been observed containing the entire text, including Internet headers, for up to seven levels of previous messages. Often, the amount of new information contained in messages posted to Usenet groups is extremely small; the rest was quotations of quotations of quotations.

48.4.3 Buying on the Web. Employers may decide to allow reasonable (however they decide the term) use of corporate resources for non–work-related activities, including buying services and products through the Internet. However, it is in the interests of employers to educate employees to avoid becoming victims of criminals. An employee distraught over the loss of significant sums due to foolish credulity will not be as productive as usual; in any case, no one wants to see friends and colleagues cheated.

Buying from known merchants through the Web can be as satisfying as buying in their stores. If you know the organizations selling goods and services, there is no more reason to be worried about buying from them through a Web connection than buying from them over the phone or in person at a store. Websites belonging to recognized merchants or organizations, such as nonprofit charities, are trustworthy, especially if they show any of several symbols representing compliance with various standards of security for customer data. Some of the safety seals in common use include SSL SiteSafe Certificates, TRUSTe, McAfee SECURE, and WhiteHat Security Certification.

48.4.3.1 Dynamic Pricing. One controversial technique that some firms have been studying is *dynamic pricing*. Dynamic pricing presents different prices to different customers. By building a profile of a specific customer's buying habits, vendors can inflate prices for people who appear to be more willing to buy higher-priced goods and lower prices for those who are cost conscious. Many brick-and-mortar stores do the same, in that stores in some parts of town may cater to richer people than in other areas; similarly, some chains of stores have been documented as charging higher prices to poor people in ghettos than in suburbs, in part because there is less competition in poor neighborhoods and the cost of doing business may be higher there. A different kind of dynamic pricing occurs in the airline industry, where seats on planes vary in price according to when they are booked and how many seats are expected to be sold. However, unlike these examples, dynamic pricing on the Web resembles traditional automobile sales, where research confirms that women and racial minorities are consistently offered higher prices than the deals for white males. In both automobile sales and dynamic pricing on the Web, the fundamental difference from the normal free-market model is that the prices are varied secretly so that only the victim of the predatory pricing sees the offered price. Without mechanisms for sharing information among purchasers, this model of pricing seems to put the buyers at an immense disadvantage with respect to the seller. It will be interesting to see how it develops over time.

48.4.3.2 Privacy. Another key area of concern when buying products on the Web is privacy. Many consumers prefer their buying habits to remain their own business. Receiving unwanted paper mail or email, because of past purchases, seems intrusive and irritating to them; they classify all such promotions as junk mail. Other consumers appreciate the convenience of receiving targeted information about new products and special sale prices for items they have previously bought. In either case, it is important to pay attention to the privacy policies offered by online vendors.

Marketers must decide whether to set up their systems on an opt-in or opt-out basis. If marketers choose the former, then all individuals actually must *agree* to have information about themselves included on lists that may be used within the organization or sold to or traded with third parties. If the system is set up for opt-out, then everyone's information may be freely disclosed, except for those who specifically state that they do not want the list keepers to do so. These are broad general outlines; the privacy policy of each organization must be spelled out in detail.

Some sites such as online bookstores and music services may keep detailed records of what each person buys from them and even what items are simply looked at. These Websites can then tailor their sales presentations to products that are appropriate to each customer's interests. Amazon.com, for example, tries to be helpful to visitors by suggesting books that may interest the returning visitor based on previous behavior. However, one of the unexpected consequences of customer profiling is that the practice may reveal more than users would wish; if you watch one of your employees enter such a Website and discover that the predominant theme is, say, weapons and techniques of terrorism, you might want to have some serious discussions with your human resources staff. A less positive application of profiling caused a flurry of interest when information about the purchasing habits of employees of specific companies was accidentally made available to those companies' competitors.

Another issue often raised in discussions of privacy involves *cookies*. Cookies are small text files that a site stores on a visitor's hard disk to store information that can be used the next time the user visits the site. Properly defined cookies can be used only by the site that deposited them. The information stored can include the sequence of Web pages the visitor saw, or personal identifiers that allow the Web software to recognize the visitor so that the Website can build up a preference profile for each visitor or client and to enable those cheery greetings like "Welcome back, Bob! We have a special deal for you on the newest title in *The Real Man's Guide to Heavy Artillery* series!" Cookies also may be used to accumulate items in a shopping cart; without cookies, each purchase would have to be concluded separately.

In general, cookies are harmless. If you do not like the idea of having identifiers stored on your system, you can block cookies in your browser settings, block them globally or on a site-by-site basis using a personal firewall, or install cookie sweepers that get rid of all cookies whenever you activate them.

For a review of legal aspects of privacy in cyberspace, see Chapter 69 in this *Handbook*.

48.4.3.3 Online Auctions. The theory behind an auction is that the competition for an object or service helps participants determine a fair price. This process can be corrupted in a real-world, physical auction if the seller conspires with confederates to bid up the price artificially. Unfortunately, this is even easier online, where anyone can have as many identities as he or she wants. The ease with which browsers and email systems allow forged headers and forged identifiers means that sellers can inflate the price of their own offerings.

The Federal Trade Commission of the United States reports that online auctions cause the largest number of complaints they receive annually about fraud.

This theoretical discussion does not even begin to address such questions as whether the auctioned items really exist, are as described, or will ever be delivered. A case of such fraud occurred on eBay, where Robert Guest of Los Angeles admitted in court in July 1999 that he defrauded victims of about $37,000 by offering goods for auction via eBay but failed to deliver anything. The customers of Mr. Guest certainly found out the

hard way that they were being cheated, but it appears that they could not have known in advance that he was untrustworthy. Although eBay maintains a system whereby potential bidders can see reviews and comments posted by earlier customers of each seller, new sellers such as Mr. Guest have no record, and anyone with a bad record can assume a new identity.

eBay has further responded to these concerns by suggesting the use of escrow services and by warning its users that it does not guarantee the legitimacy of the transactions it facilitates.

There are also concerns about the legality of some of the items put up for auction. Someone offered items made from endangered species, in violation of the Convention on International Traffic in Endangered Species (CITES). The products included dried feet of elephants and gorillas caught in snares and allowed to die excruciating deaths before being hacked into pieces. In the United States, buying, selling, and possessing such contraband can lead to arrest, prosecution, fines, or imprisonment.

More ludicrously, someone put up a human kidney for sale through eBay in September 1999 and received bids of up to $5.8 million. The auction service canceled the sale because selling human organs is a federal felony punishable by up to $250,000 in fines and at least five years in jail. A week later eBay had to shut down an auction for an unborn human baby. Prices for the supposed baby had risen into the $100,000 range before eBay shut down that auction. Finally, a fool or a prankster—it is unclear which—tried to sell 500 pounds of fresh marijuana online. The auction was shut down after 21 hours, during which prices offered had reached $10 million. In August 2001, a couple offered to name their baby in accordance with the wishes of a high bidder. That auction, too, was ended prematurely.

Most of the bids probably were not legitimate. It is unlikely that everyone who bid for kidneys, pot, and babies really expected to pay for what they were bidding on. They may have been treating the auction like a video game, with no element of reality. Situations such as these invite other abuses, and ordinary users are often at a loss as to how to proceed.

Even if the items being offered for sale online are ordinary things such as software or physical products, they may have been obtained illegally. Online auctions are a frequently used channel for fencing stolen goods.

Corporate users should probably not be using Internet auctions to buy or sell products, except in those closely guarded, industry-specific sites that have proven their worth. Certainly, employees should not be using corporate Internet access to engage in such activities for their private purposes.

48.4.4 Online Gambling. It is hard to imagine that any enterprise would authorize employees to gamble online using corporate resources, but providing employees with the following guidance may be a valuable service.

48.4.4.1 Fraud and Error. In 1998, the Arizona lottery discovered that no winning number in its Pick 3 game had ever included even one numeral 9.[9] It turned out that the pseudorandom number generator algorithm had an elementary programming error that generated only the digits 0 through 8. All those who had used a 9 in their lottery numbers felt justifiable anger—especially when they were told they could have a refund, but only if they had kept their old losing tickets.

The Arizona lottery used a simulated random process to provide the illusion to gamblers that they were betting on a physical process such as balls mixing together in a barrel and falling out of a tube. One of the problems with the Arizona simulation is similar to a genuine vulnerability in proprietary (i.e., secret) cryptographic algorithms.

As cryptographers have stressed over many decades, the security of an encryption scheme should not depend on the secrecy of its algorithm. Had the lottery algorithm been exposed to public scrutiny, its flaws would have been detected sooner. For example, in the 1980s, there was much excitement over a new encryption scheme called the *knapsack algorithm*; after extensive examination by cryptographers, it proved to be flawed. It is conceivable that someone detecting the flaw in the Arizona lottery might have made bets with a higher probability of winning than those of uninformed people, but exposing the algorithm and its implementation to scrutiny before it went into production would have made that less likely.

These examples demonstrate that electronic gambling, as in older, conventional types, is subject to more than the rules of chance. Lack of conformity to good security practices lays both the gambler and the house open to abuse and to inadvertent errors.

48.4.4.2 Lack of Control. Physical gaming devices are located in real-world establishments under the nominal control of regulatory and law enforcement officials. Even so, they are always adjusted for a certain predetermined payout. Gambling based on the results of actual sports events or contests is validated by external news reports, although the contests themselves can be rigged. But there is no basis for a gambler to trust the results of computer-generated pseudorandom numbers displayed on a browser screen.

Most individual gamblers will never know if a long-range analysis of the pseudo-random numbers would support their hopes for fairness in the odds. No one is keeping track of these data except the people making money from the participants, and they are not distributing the results.

The disclaimer at one Internet gambling portal, findinternetcasino.com, is not very encouraging:

> Although every attempt has been made to ensure fairness and security toward the player at each of the links that can be found in the directories, FindInternetCASINO® cannot be held responsible if discrepancies occur between an Online Gambling operation and you, the player, after following a link from this WWW site. Consult your local authorities prior to registering with any online wagering service. U.S. Citizens: The information at this site is for entertainment and news purposes only. Use of this information in violation of any federal, state, or local laws is prohibited.

48.4.4.3 Legal Issues. In some jurisdictions, betting online is illegal. In the United States, for example, it is already illegal to use interstate telecommunications to place bets; in addition, Internet betting is illegal in the United States even if the host is outside the United States. At the same time, due to ambiguities in the current laws and the inability to clearly enforce them, the use of overseas betting sites has driven this business to a total of over $15.5 billion a year, over half of that income coming from the United States. The ambiguities stem from a lack of a clear definition on what constitutes illegal online gambling. This resulted in certain groups of individuals believing that they were exempt from the law, poker players being the most common. In addition, online horse racing receives a specific exemption from the law, but without accompanying clarification on whether the wagering process constitutes online gambling.

Unfortunately, in the United Kingdom and many other countries, online gambling is for the most part legal. This creates numerous conflicts of interest, and international tension, between various betting companies in legalized countries, all advertising to Americans eager to risk their money for a chance at a big payout. It appears that until clear definitions are included in the law, the blurred line between legal and illegal gambling activities using online resources will continue.[10]

48.4.5 Internet Addiction. Any activity can become the basis of compulsive exaggeration. A small proportion, around 5 percent, of Internet users may qualify as addicted to any of these computer-mediated activities:

- An uncontrollable desire to find and organize more and more information about an enormous range of topics
- Excessive involvement in games, gambling, and buying things on the Internet
- Excessive concentration on relationships mediated through email and chat rooms, to the detriment of real-life relationships
- Involvement in long sessions of viewing pornography, or of being sexually stimulated via email, chat rooms, pornographic sites, or sexual-fantasy games

None of these activities is a suitable use of corporate computing resources, and employees should be alerted to the policies prohibiting such activities at work. In addition, everyone should be aware of the dangers of Internet addiction.

The issue here is what constitutes *excessive* involvement in these activities. Professional psychologists such as Dr. Kimberly Young have identified some of the diagnostic criteria for these disorders, including these based on her *Internet Addiction Test*[11]:

- Regularly staying online longer than intended
- Often neglecting obligations to spend more time online
- Consistently preferring to spend time online instead of with one's partner
- Frequent complaints by friends and family about excessive Internet use
- Suffering consequences at school or at work because of time spent online
- Giving email a higher priority than other important issues
- Concealing the extent of Internet usage
- Turning to the Internet as a substitute for dealing with disturbing issues
- Feeling that life without the Internet would be devoid of meaning and pleasure
- Getting angry when disturbed during Internet usage
- Losing sleep due to late-night Internet activity
- Yearning to be back online

Those who feel uncomfortable about their level of involvement with the Internet would do well to take this test offered by Dr. Young, and, if several of their answers are positive, to seek counseling to prevent possibly tragic consequences of untreated addiction.

48.4.6 Online Dating and Cybersex. As in other topics in this chapter, it is unlikely that corporate policy would allow users to engage in online dating and cybersex. Nonetheless, in line with the overall orientation of this chapter, the next sections will help employees understand the issues in these online activities.

48.4.6.1 Dating Online. Thousands of sites on the Web specialize in helping people meet each other. In a sense, chat rooms and bulletin board systems are ways for people with similar interests to communicate about their hobbies and lifestyles. There are also sites that specialize in helping people find others who match particular profiles.

Some of these sites are free; others charge fees for participation. Dating service sites usually explicitly restrict participation to people over 18 years old, and most of them depend on possession of a credit card as their sole mechanism for authenticating age. It is very difficult to exclude teenagers, or even younger children, from such sites if they have access to credit card numbers.

Parents, teachers, and employers who want to get a sense of what is going on can type "online dating" in the search field of a search engine such as Google (www.google.com) and then visit a few of the sites. If children post information about themselves in such a cyberspace locale, even with false information claiming that they are adults, there is a real risk of attracting unsavory characters or perhaps ordinary people who can become angry at being tricked into exposing their feelings to an imposter.

48.4.6.2 Sex Talk Online. In addition to matchmaking, users of the Internet also can get involved in *cybersex.* People chatting online can describe themselves or each other in sexual interactions that are inappropriate for youngsters. Such online chat also has been implicated in a number of divorces, since many spouses find it wholly inappropriate that their beloved is getting sexually excited with a stranger via the Internet.

In August 2001, a 15-year-old girl from Massachusetts was alleged to have been kept captive for at least a week during which she was repeatedly sexually abused by the couple who had brought her to Long Island. According to the criminal complaint, she was also loaned out for two days to another man and further abused. The couple had met the teenager in an Internet chat room, where their conversation was explicitly sexual.

In the work environment, circulating sexually charged messages or outright pornography can justifiably be perceived and described as fostering a *hostile work environment* and can lead to lawsuits by the affected employees.

Employers should promulgate policies to prevent such abuse and monitor corporate email and instant messaging to ensure that no one in their employ engage in these activities using corporate resources.

48.4.6.3 Traffic in Women. A number of sites on the Web, particularly some situated in the former Soviet bloc, advertise services for introducing men to willing candidates for marriage. The evidence is strong that much of the information communicated about the supposedly nubile and marriage-oriented women is false. Many of the pictures are taken from public Websites and include actresses and people who have posted their photos on social networking groups. Sometimes the same picture has dozens of names associated with it. Much as in the phone-based sex-talk services, people claiming to be youthful, attractive, persons of marriageable age may be nothing of the sort, and may be copy/pasting responses from prepared scripts. When men travel to visit their potential mates, they can be charged high rates for the privilege of taking their dates to expensive restaurants. Some of the women who actually do go through with marriages later divorce their hapless victims once they are admitted to their husband's country of residence in what appears to be systematic fraud.

48.4.7 Games and Virtual Reality. Some enterprises allow their employees to play games at various times during the day—usually low-usage times such as lunch, or before and after the normal workday. However, some Internet-enabled multiuser games can consume enormous bandwidth; the shoot-'em-up (*first-person shooter*, or FPS) game called Quake was notorious in its day for saturating all available connectivity.

When helping employees understand how to negotiate the perils of the Internet, you might recommend that parents read reviews of video games before allowing their young children to play them. Some games have astonishing levels of graphic violence ("Brilliant Bleeding! Detailed Decapitations!") and unusual values ("Win points by burning as many residents to death as possible!"). This latter example is based on a notorious case in which a video-game vendor was apparently surprised by the public wave of revulsion over a game that glorified arson. Some military and police shoot-'em-up games explicitly take points off for hitting innocent bystanders; others do not. Some games use graphic nudity; others are more modest. The main point is that relying on the judgment of eight-year-olds to choose their own entertainment may be unwise.

From a corporate perspective, it would be unusual to find employers encouraging the use of local or networked games during working hours; however, some may allow use of their resources in off-hours, assuming the corporation does not maintain around-the-clock operations. However, issues of suitability persist; some games may contribute to a hostile work environment and lead to complaints and lawsuits from offended employees.

A development that started in the 1990s has become a potentially valuable tool in the first decades of the twenty-first century: virtual reality or virtual worlds, such as Second Life (http://secondlife.com). These services use controllable representations called *avatars,* which allow some degree of expressiveness when communicating. Participants see a representation of a three-dimensional world, complete with viewpoint and perspective, that includes their interlocutors in a shared virtual reality that can be creative and fun. Some companies are using resources in these virtual worlds for advertising, delivery of services (e.g., training and education), and internal remote meetings or training. Organizations must determine appropriate policies about the use of such services.

48.4.8 Changing Email Addresses. Employees may leave a company, change organizational units, or change their own names. All of these changes may result in new email addresses. Handling such changes poorly can lead to trouble.

One reaction to such a change is to delete the original email address without notification to anyone. Email sent to the original address is returned with an *undeliverable (no such user)* error. The sender must then find out what happened—or may simply drop the connection altogether, possibly losing a company a client or leaving important information undelivered.

If the disappearing email address is due to the user's change of name (for example, resulting from marriage or divorce), it is even possible that the corporate directory will have the original name wiped, making it difficult for correspondents unaware of the new name to reach the person at all by email or by phone.

A wiser response to any such change is to autoforward the incoming mail to the correct address. For example, if Farid Hallings's original email address was fhallings@company.com and is now fmalteso@company.com, any mail sent to the first address would automatically end up in the mailbox for the second address. If Farid no longer works for the company at all, the email can be forwarded to the appropriate replacement's address. Such forwarding can be maintained for whatever period seems appropriate.

In addition to the autoforward, it may be helpful to send automatic notifications to the sender of the outdated email. "Farid Halling's new email address is fmalteso@company.com; your message has automatically been forwarded. Please make a note of the change in your email address book."

48.5 LEGAL LIABILITY. This section briefly reviews some of the legal issues that may arise as a result of misuse of email and Internet resources. For more detailed information, see Chapters 63, 64, 69, 70, 71, and 72 in this *Handbook*.

48.5.1 Libel. Some people have taken advantage of the freedom to publish whatever they want by crossing the boundaries of libel. For example, the self-styled reporter Matt Drudge went too far in postings on his electronic scandal sheet in 1997, when he made unsubstantiated accusations about White House advisor Sidney Blumenthal's marriage. Professional journalists pounced on him for shoddy journalism. Blumenthal and his wife filed a $30 million libel suit against Drudge even after he apologized for failing to verify the gossip he disseminated. Drudge then claimed that public White House support for Blumenthal amounted to a threat against free speech.

In another notorious case, Walter Cronkite, whom polls revealed to be the most respected man in the United States in the 1980s, was appalled to discover a page of lies about him on the Web in 1997. A 28-year-old programmer, Tim Hughes, invented and posted a scurrilous story about Cronkite's becoming enraged at the author, shrieking imprecations at Hughes and his wife, boasting about his own infidelity, and spitting in their spice cake at a Florida restaurant. In addition, the anti-Cronkite Web page included falsified photographs purporting to show Cronkite at a Ku Klux Klan meeting. Cronkite threatened to sue for libel; Hughes took the page down and weakly protested that it was all a joke.

The effect of this kind of misinformation on children or immature employees, untrained in critical thinking and lacking in skepticism about information on the Internet, can be damaging.

Another source of information is the Usenet—that collection of thousands of discussion groups on every conceivable topic. These discussion groups fall into two major classes: moderated and unmoderated. In a moderated group, messages are passed through a *moderator* who decides either to post them for participants or to delete offensive or otherwise inappropriate messages. Not all moderated groups are reliable, and not all unmoderated groups are unreliable. However, many unmoderated groups distribute unsubstantiated information from people who appear to derive their major pleasure in life by insulting other participants and by making outrageous statements about any topic that comes up. Everyone should be trained to recognize emotional and inflammatory language, and should be encouraged to apply skeptical analysis to all statements, especially to those published in rants.

In the first decades of the twenty-first century, *blogs*—commentaries published on the Web by individuals or groups—have exploded into common awareness. The same principles of critical evaluation apply to blogs as to any other source of disintermediated information. In one case involving author Kabay, a lunatic conspiracy site illegally reposted the entire text of one of his columns describing InfraGard with pictures of victims of Nazi atrocities in concentration camps between each paragraph. A Digital Millennium Copyright Act (DMCA) *takedown request* was obeyed, but the incident was disgusting to the author.

48.5.2 Stolen Software, Music, and Videos. Organizations cannot permit employees to download and make illegal copies of intellectual property of any kind. Security policies must explicitly address these issues; security monitoring must explicitly control for abuse of corporate resources in such activities. The risks to organizations by tolerating such violations of law are severe. For more details of intellectual property law, see Chapter 11 in this *Handbook*.

48.5.3 Plagiarism. A different kind of fraud involving intellectual property occurs when people misrepresent someone else's work as their own. Older students know intellectually that this is supposed to be bad, but for young children, the issue is completely abstract. The problem today is that plagiarism is easier than ever and harder for teachers to detect.

Academic guidelines try to make it clear to students that copying other people's work without attribution is called *plagiarism* and is severely frowned on. Plagiarism includes not only direct quotation without indications of origin but also paraphrasing that merely shuffles the ideas around a little or substitutes synonyms for the original words. In many institutions, plagiarism is grounds for suspension or expulsion. In all cases, plagiarism defeats the purpose of writing assignments by eliminating the opportunity for critical thinking and creative expression. Few plagiarists remember what they have copied from others after they hand their material in.

Assuredly, students have traded term papers and other assignments for centuries. However, the availability of electronic documents and of the World Wide Web has enormously increased both the fund of material that can be plagiarized and the ease of copying. Worse still, some people are profiting from easy accessibility by selling papers specifically for plagiarism and even writing papers to order. In one study by Peggy Bates and Margaret Fain of the Kimbel Library at Coastal Carolina University, the authors easily located over 100 sites on the Web selling or donating papers to students for plagiarism.[12]

To combat this problem, science has come to the aid of beleaguered instructors by providing automated similarity analysis of any paper submitted electronically. The system uses a bank of more than 100,000 term papers and essays as well as documents located on the Web; analysis uses pattern recognition to measure similarities among different documents and to estimate the probability of plagiarism. According to the turnitin.com documentation:

> Our system is now being used in the majority of universities in the United States and the U.K., as well as a large number of schools around the world. Many of these institutions, among them UC Berkeley and the fifty-eight member schools of the Consortium of Liberal Arts Colleges, an association of the most respected liberal arts schools in the US, have chosen to ensure the academic integrity of all their students by selecting institution-wide subscriptions to our service. Other universities, such as Harvard and Cornell, have elected to make use of our system on a departmental or single-instructor basis.

Plagiarism is also a risk to the enterprise; having employees misuse other people's or other organization's materials without attribution can lead to lawsuits, embarrassing publicity, and serious financial penalties. In one notorious case from 2003, a policy paper about Iraqi intelligence organizations distributed by the Prime Minister's Office in the United Kingdom was discovered to include large swathes of verbatim material, including typographical errors, copied and pasted without quotation marks and without indication of its source.[13]

Practical guidelines:

- Discuss plagiarism clearly at work, at home, and at school.
- Use examples to illustrate the difference between plagiarism and a legitimate use of other people's work.
- Encourage children to practice summarizing information in their own words.
- Practice writing references to quoted material.

- Have a student submit their term paper to one of the online services that verify originality.
- Discuss how antiplagiarism sites analyze documents to measure similarities and help teachers identify plagiarism.

48.5.4 Criminal Hacking and Hacktivism. As discussed in Chapter 45 in this *Handbook*, it is important that all employees understand and agree that using corporate systems for unauthorized access to computers and networks is grounds for dismissal, and possibly criminal prosecution. In particular, no employee should ever imagine that testing for security weaknesses in the enterprise's systems without authorization is a contribution to security.

The motivation for illegal actions does not mitigate the seriousness of computer trespass. Employees should be informed explicitly that regardless of the excuse, no violations of law will be tolerated. For example, hacking into systems in another country to support a war effort is not excusable; nor is destroying child pornography sites a good idea. Cybervigilantes can destroy evidence needed for prosecution.

48.5.5 Creating a Hostile Work Environment. In today's society, there are numerous activities and language constructs that individuals of a certain race, gender, sexual orientation, national origin, religious affiliation, or other legally protected characteristics may find offensive. Any type of harassment, most especially comments or actions based on these protected characteristics, toward another employee may create a hostile work environment. The two most common situations created by a hostile work environment are:

1. A reduction or loss of productivity due to the harassment, whether physical, verbal, or psychological
2. A reduction in salary, bonus, job level, responsibilities, or other components of compensation due to one or more of the legally protected characteristics

Although there are no formal laws barring hostile work environments, Title VII of the Civil Rights Act of 1964 covers these types of situations. These laws are written in such a way that an individual comment or action does not usually constitute harassment. Rather, a pattern of frequent, severe, and pervasive abuse may constitute a hostile work environment. It is important to distinguish between *quid pro quo* harassment, where an employee is required to tolerate such harassment in order to maintain job status or compensation levels, and a hostile work environment. Both are very serious and potentially illegal activities, but this section focuses on the hostile environment.[14]

Employers are obligated by law to set appropriate expectations around employee behavior, and confidentially and swiftly to investigate any complaint of harassment from an employee. Employees are granted some legal protections such that retaliation by an employer for sounding the alarm on a hostile environment is illegal.

These issues become even more important when office romances occur. Although many employers forbid couples from working together in the same department, in part to avoid any perception of favoritism or future harassment cases, should the romance fail, there is still the potential for romantically linked coworkers to create a hostile environment for others. In this case, employers have an obligation to define in policy that coworkers should maintain a professional relationship while on company business. Even though two individuals may feel that their words or actions are seemingly

innocuous, it is the perception of others around them that creates the basis for a harassment complaint. The best solution is to keep personal lives out of the office, which is a difficult but appropriate recommendation for everyone.

48.5.5.1 Hate Groups. Another source of concern for employers and parents is the easy accessibility of hate literature on the Web. Hatemongers have taken full advantage of the largely unregulated nature of the 'Net to spread their pernicious messages. One can find Websites devoted to hatred of every imaginable identifiable group. Race, ethnicity, religion, gender, sexual orientation, immigration status, and political ideology—anything can spark hatred in susceptible personalities. Unfortunately, some of the hate groups have been quite successful in recruiting young people through the Web; they publish propaganda such as pro-Nazi revisionist history that may fool uncritical people into believing their rants. Neo-Nazi and racist skinhead groups have formed hate-rock groups that take advantage of kids' enthusiasm for very loud music with aggressive lyrics.

Employers cannot tolerate the slightest involvement of their employees in such activities using corporate resources. Aside from their possible personal revulsion at such hatemongering, managers also should be aware that toleration of intolerance can lead to a hostile work environment in which targets of hate or contempt can legitimately appeal to the courts for compensatory and punitive damages. Employees must understand and agree that using any corporate resources for participation in hate groups is a serious infraction of Internet usage policy.

According to the Simon Wiesenthal Center, there are over 2,300 Websites advocating hatred, of which over 500 are extremist sites hosted on American servers but authored by Europeans; most European countries have strict antihate laws. Using more stringent criteria, the Hate Watch group estimates more than 500 extremist hate sites on the Web; it distinguishes between hate propaganda and those pages that consist largely of racial epithets, dismissed as mere graffiti.

The Southern Poverty Law Center monitors 500 active hate organizations in the United States. It has regularly reported on the growing number and stridency of such sites. In comments about the center's paper for the United Nations Commission on Human Rights, spokesperson Mark Potok said at a conference in 2000:

> A few years ago, a Klansman needed to put out substantial effort and money to produce and distribute a shoddy pamphlet that might reach a few hundred people. Today, with a $500 computer and negligible other costs, that same Klansman can put up a slickly produced Web site with a potential audience in the millions.[15]

A fundamental reality is that human beings are gregarious. They find it very easy to affiliate with others to form *in-groups,* groups to which they feel entitled to belong. Unfortunately, defining in-groups naturally means it is equally easy to define *out-groups*: groups to which we *do not* want to belong. Grade school and high school cliques are examples of in- and out-groups. A wealth of study in social psychology confirms the validity of the universal impression that we tend to inflate our esteem for in-groups and to reduce our respect and liking for out-groups. However, research also shows that social norms against discrimination can reduce hostility toward out-groups; thus it seems likely that parental and teacher articulation of norms of tolerance can significantly reduce children's susceptibility to the blandishments of hate groups.

48.5.5.2 Pornography. Pornography—even with the most restrictive definitions—is widespread on the Internet. Observers of 'Net culture have commented

that the sure-fire way of telling if new technology is going to be a success on the Internet is to see how quickly pornographers can apply it. For example, the appearance in July 2000 of the first WAP (wireless application protocol) pornography sites signaled the adoption of WAP technology into the mainstream. Although the sites offered only tiny grainy images of naked Japanese models, sociologists said that the same expected sequence of rapid technological advances had occurred with photography and video cameras.

48.5.5.2.1 Prevalence of Porn. Some studies of Internet traffic have claimed that more than half of the total 'Net bandwidth is used for transfer of pornography or solicitations for purchase of pornography.

48.5.5.2.2 Trickery. Pornographers use various tricks to get people onto their Websites:

- Using a different domain, like the old whitehouse.com, which used to take advantage of interest in "whitehouse.gov" by showing porn (it is now a directory with several paid links for dating).
- Misspellings, such as the now-inactive micosoft.com, which traded on the likelihood of mistyping "Microsoft.com."
- Junk email invitations with innocent-looking labels for URLs that do not match the actual link but instead take the viewer to a pornography site.
- Padding porn-site metatags (normally invisible text used to describe a Website) with inoffensive keywords that place the site high on search engine lists where they can appeal to children.
- Disabling normal features of a browser to trap victims in the porn site. One perpetrator who was shut down by the Federal Trade Commission (FTC) actually ran Java applets that disabled the back arrow and defeated the ability to close the browsers. People trapped in porno-hell had to reboot their computers to get out.

Porn sites are notorious for using deceit to defraud their victims. One widely used scam is to demand a credit card number from a visitor as proof of their age (it is nothing of the sort), then to charge the card even though the site clearly states that there is a period of free use.

In 1996, viewers of pornographic pictures on the sexygirls.com site were in for a surprise when they got their next phone bills. Victims who downloaded a *special viewer* were actually installing a Trojan horse program that silently disconnected their connection to their normal ISP and reconnected them (with the modem speaker turned off) to a number in Moldova in central Europe. The long-distance charges then ratcheted up until the user disconnected the session—sometimes hours later, even when the victims switched to other, perhaps less prurient, sites. Some victims who stayed online for a long time paid more than $1,000 in long-distance charges. In February 1997 in New York City, a federal judge ordered the scam shut down. An interesting note is that AT&T staff spotted the scam because of unusually high volume of traffic to Moldova, not usually a destination for many U.S. phone calls. In November 1997, the FTC won $2.74 million from the Moldovan telephone company to refund to the cheated customers—or the ones willing to admit to having been cheated.

Both of the scams just described relied in part on the reluctance of porn-seeking victims to admit to their socially disapproved interest. Few victims were willing to pursue the matter until the damages mounted into the thousands of dollars.

48.5.5.2.3 Filtering. An entire industry has grown up to try to shield (or block) children from seeing pornography or other materials deemed offensive by their parents or by the makers of the blocking software. The popular blocking systems are reviled by many free-speech advocates, and often ridiculed for what are described as clumsy, keyword-oriented algorithms. The classic examples of ludicrous blocking include trapping access to any site that uses the word *breast*—including even possibly this very page if you are reading it on the Web. Other simple-minded traps have blocked users from accessing information pages for geographical locations ending in the old British suffix *-sex* such as Wessex, Sussex, Middlesex, and so on. The village of Scunthorpe in England was blocked by software used by a major Internet service provider because its internal filters prevented anyone from using vulgar words in their mailing address.

Some of the blocking software products use hidden assumptions about the unsuitability of a wide range of topics, including abortion rights, civil rights, political ideology, and gay liberation. Any parent is entitled to express opinions about any topic; however, parents will want to check on whether a particular program is imposing its makers' political agenda by stealth. In the workplace, employers who use broad-spectrum blocking software may interfere with legitimate research by their employees.

48.5.5.2.4 Monitoring. A different approach to interfering with the nefarious deeds of pornographers is to install monitoring software on the computers that employees use at work or that children will use at home. These products keep a log, or audit trail, that allows employers and parents to see exactly what users have been doing with their computers.

In the family context, most important, however, is the principle that machines and programs cannot by themselves teach values. Instead of relying only on passive barriers or on snoopware, parents would do well to make surfing the Internet a family activity rather than a private hobby. When kids express interest in pornography—because our popular culture is full of sexual innuendo that children read, hear, and see—it makes sense to discuss the issues rather than try to pretend that they do not exist. One approach for reducing the power of the forbidden fruit offered by pornographers is to explain to children in a supportive and nonpunitive way why sexual exploitation and degradation are bad for people. Children who stumble on porn sites by accident or at their friends' houses may be better prepared to cope with the sometimes disturbing images and words if their parents have prepared them for this aspect of today's world.

48.5.6 Archiving Email. Organizations must remember that email may be demanded as evidence in court cases. There is a fiduciary duty to maintain business records appropriately for each type of business, and that obligation extends to electronic records. Policies should stipulate how long email records should be maintained. Destruction of email should never be selective, especially if there is an anticipated threat of legal action. Selective destruction of particular records, or premature wholesale destruction of email, may be interpreted by the courts as grounds for charges of interference with the judicial process.

For details of backup and archiving policies, see Chapter 57 in this *Handbook*.

48.6 RECOMMENDATIONS. This section summarizes some practical recommendations for employees and their families. Framing policies in a way that supports employees' concern about their own families is a helpful way of increasing the perceived value of the guidelines.

48.6.1 Protecting Children

- Explain the dangers of communicating with strangers via the 'Net in the same terms that you discuss the dangers of talking to strangers anywhere else.
- Alert children to the questionable identity of anyone they meet exclusively through the 'Net or via email. Discuss the possibility that people are not what they claim to be in their online persona.
- It is important that children feel confident of a supportive response from their parents when raising these issues. Establish a calm atmosphere so that children will not fear your reactions if they are troubled by what they encounter online. Worst of all would be to punish a child for reporting a disturbing incident.
- Tell children not to give their address to strangers they meet electronically.
- Children should not send pictures of themselves to strangers.
- Make a practice of discussing online relationships in a friendly and open way at home. Show interest in the new friends without expressing hostility or suspicion; ask to participate in some of the online chats and email correspondence. Invite your children to sit in with you during your own online interactions.
- If a child feels that another child met online is becoming a good friend, parents should contact the child's parents by phone and, eventually, in person before allowing contacts.
- If a child wants to meet someone encountered on the Internet, be sure that a parent is involved at all stages. Never let a child meet anyone in the real world whom he or she has met only on the 'Net. Any attempt to induce a child to meet the correspondent alone or secretly should be reported to local police authorities for investigation.
- Make it clear that anyone who suggests hiding an online relationship from the child's parents is already doing something wrong.
- Make it clear to your children that no one has the right to send them age-inappropriate, sexually suggestive, or frankly pornographic materials, whether written or pictorial. Suggestions on the Internet that children engage in virtual sex play or sexual fantasies should be reported to parents right away. Making, transmitting, and storing child pornography is a felony; report such cases to local police authorities at once.
- Children receiving a request for anything unusual (e.g., a request for a piece of clothing or for nude pictures) should immediately report the incident to their parents. Teachers and other caregivers can adapt these principles for the specific circumstances of their relationship with the children they are taking care of.

48.6.2 Threats

- Employers, parents, and teachers should clearly enunciate policies preventing anyone—including children—from uttering threats of violence or other harm, even in email messages or chat rooms.
- Employees should be instructed to report all threats directed at them, or at others, to the security officers in their organization; similarly, parents, teachers, or librarians should ensure that children know to report any threats immediately to the appropriate adult.

48.6.3 Hate Sites

- To protect children against the wiles of these hateful people, the most important step is to discuss the issue of hate speech and hate groups with them openly. Parents may even want to visit some of the sites listed below *with your kids* to give them a sense of the problem and possible countermeasures.
- Discuss your children's feelings about out-groups in their own lives; for example, encourage them to speak freely, without fear of punishment or reprimand, about whatever groups they do not like. Then pursue the discussion with explanations of such issues as cultural differences, history, or whatever else you feel will help your children gain perspective on their own feelings and behavior. Of course, this positive attitude cannot be applied to hate groups or similar outlaws.
- Provide positive social role models for children with respect to hate groups. Speak out firmly in opposition to intolerance rather than sit silently by when bigots display their hatred for other groups.

48.6.4 Pornography

- Place young children's Internet-access computers in a family area of the home rather than in their bedrooms.
- Interact with your children while they are using the Internet; treat the Web browser like a window on the world, and be present to help your children interpret that world in a way consistent with your values.
- Talk with your children about the existence and nature of pornography; as they reach puberty, assure them that there is nothing wrong with being interested in sex, but that pornography is not a healthy way of learning about wholesome, loving relations.
- Warn your children about some of the tricks used by pornographers to get traffic on their Websites, such as telling them to download special readers. Tell them about the Moldovan porn scam.
- Discuss the issue of junk email that advertises porn sites. Warn children that no one should ever click on a URL from any kind of junk email because it can easily be a trick to get them into dangerous territory.
- Teach your children to keep an eye on the actual URL that appears in the browser window; any discrepancy between the visible URL shown on a page and the actual URL should alert them to the possibility of fraud.
- Explain that pornographers sometimes charge for access to their sites without permission; be sure your children understand how dangerous it would be to give your credit card number to these people for any reason.

48.6.5 Internet Addiction

- Know the warning signs of Internet addiction and self-monitor.
- Discuss Internet addiction and its warning signs with your employees and your children.
- Encourage open discussion of feelings about the 'Net, so that children feel free to turn to you for help if they become uncomfortable or unhappy about their own experiences on the 'Net.

48.6.6 Online Dating

- Do not build online profiles or give out addresses, phone numbers, or school names.
- Do share email accounts with your children, and oversee their messages.
- Keep the computer in a family room where children's activities can be monitored.
- Remember that people may lie when describing themselves online.
- Do not allow children to meet online users without permission, and make all meetings in public places with adult supervision.
- Forward copies of suggestive or obscene messages to your Internet service provider.
- Find ways to block objectionable material.
- Discuss online dating with kids so they understand what is involved.
- Ensure that kids understand why it is inappropriate and even dangerous for them to masquerade as adults in online dating services.
- Do not rush into face-to-face contact; you need to be sure that you are meeting someone who is on the level, not an imposter who has ulterior motives.
- You may want to take advantage of anonymizing services offered by some dating sites to avoid handing out your real email address to complete strangers.
- Be suspicious of anyone who tries to pressure you in any way, including demanding money or insisting on a meeting, before you feel confident of the person's good intentions.
- As you are getting to know someone online, ask questions about lots of things you are interested in—for example, hobbies, politics, religion, education, birth date, family background, and marital history and status.
- Keep the answers you receive and beware of people who provide inconsistent or contradictory information as they are communicating with you—any lie is a danger signal.
- Be suspicious of anyone who seems to be too good to be true; if someone matches you on every single preference or interest you mention, try mentioning the very opposite of what you said earlier in the communications and see if the person agrees with *that* too. Trying too hard to please by lying may mark a manipulative and potentially dangerous personality.
- Be honest about yourself; state your own interests and characteristics fairly, including things you think might be less attractive than stereotypes and cultural norms dictate. A mature, good person will not necessarily be turned off if you do not look like a movie star, or if you do not play four musical instruments perfectly, or if you lisp.
- If you get to the point of exchanging pictures, be sure that you see the person in a wide variety of situations and with other people; some online daters send false pictures to misrepresent themselves.
- Talk to the person you are getting interested in over the phone; be suspicious if the person resists such a request for a long time or always has excuses for not being available when you have agreed to talk.
- Listen carefully to how the person sounds on the phone, and be suspicious if you now receive information that contradicts something the person wrote to you about. Any lie should alert you to potential problems.

- Before you agree to meet, get your date's full name, address, and telephone number. Be suspicious if the person refuses to give you a home number: Could he or her have a spouse or a current live-in friend that he or she is trying to deceive? Call the home number a couple of times to see if someone else answers.
- Give the person's information, and the exact details of where and when you are going to meet, to friends and family. Do not ever accept a date with someone who wants to keep the location and time a secret. Be sure the meeting place is well lighted and in a public place such as a coffee shop.
- Do not allow a stranger to pick you up at your house, and be sure you can get home by yourself.
- Before considering further involvement, for safety's sake think about having a background check done on the person you like, using a professional service.

48.6.7 Online Games

- Learn to play some of the games your kids are enthusiastic about. Take the time to immerse yourself in the imaginary worlds they play in, and study the underlying values that are being communicated by the game creators.
- Use published reviews from online or other media that reflect your own family's values before allowing games into your home.
- Accompany your children to the stores when buying video games. Check for parental warning labels. Talk to the salespeople if you think they are reliable.
- Know the characteristics of your hardware and software before buying recently released games. Do not buy a new game only to discover that it does not run on your obsolescent system. A disappointed child can apply intense pressure to spend money on a new system. Some games are computationally intensive and require expensive, advanced computer hardware and modern sound systems, complete with a high-powered amplifier driving woofers and subwoofers.
- Try making game playing an opportunity for family fun or parent–child bonding instead of the isolating experience games can sometimes be. See if you can all have fun with puzzle- and exploration-oriented games such as *Myst* and *Riven*, neither of which involves violence, and both of which are visually beautiful.

48.6.8 Online Purchases

- Before spending a considerable amount of money on a new online merchant's site, do some basic research into the site's reliability. Check the company's reputation; see if it belongs to the Better Business Bureau (BBB), and contact the appropriate chapter of the BBB to see if there have been complaints about the vendor.
- Do a Web search using a good search engine, such as Google, to see if there are any up-to-date reports about customer experience on the site you are interested in.
- Pretend that you already have a problem and look for the customer service pages. Are there clear instructions on how to communicate problems? Would you have the choice of email, letters, or phone communications? If you have the time, you may even want to try calling customer service and find out just how they handle calls. If you hit a company that hangs up on you when its lines are busy ("We are sorry, but all our agents are busy; please call back later."), you might want to give serious thought as to whether it is safe doing business with them.

- Read the company's return policy; how does it handle breakage in transit, or defective goods? Does it offer guarantees on delivery time? What happens if the company is out of stock on a specific item—does it ship partial shipments or wait for everything to be ready? When out of stock, does it charge your credit card immediately, or only after the shipment is made? If it splits your shipment, does it charge extra for delivery of the later parts?
- Read the site's privacy policy. If the text is practically invisible 6-point yellow on white, be suspicious. Look for weasel-words in the clauses that say, for instance, that their policies can be changed at any time without notice. You must check the site regularly to see if the policy has changed, but this is unrealistic. Instead, look for firm, clear assurances that your personal information will not be sold, traded, or given away without your permission. Usually, Website owners state that they may have to divulge information to partnering organizations that handle such normal functions as billing and order fulfillment. There can be little objection to this provided the partners are bound by acceptable security policies.
- Keep a detailed record of your transactions. Use the browser functions to save copies of, or print out, the relevant Web pages with descriptions of the products, prices, a summary of your order, the order number, promised delivery date, and method of shipment.

48.6.9 Online Auctions

- Before becoming involved with online auctions, research the value of goods you are interested in buying. Check bricks-and-mortar stores, online retail outlets, and comparative shopping sites that provide you with specific prices.
- Examine the policies and costs on shipping, warrantees, and refunds.
- Set your upper limit before you get involved in an auction. Do not be influenced by the value other people appear to place on a particular product or service, and certainly do not be caught up in a bidding frenzy.
- Do not treat online auctions as a competition you have to win.
- Look for auction services that provide a guarantee of support if you are cheated in a transaction. For example, check for language in the terms of service that covers losses up to a suitable limit. Check for insurance policies, costs, terms, and limits. Use search engines to evaluate the trustworthiness of the service you are thinking of using.
- If possible, use a service that provides an escrow function so that you pay money to the service and then release it only when the product is received in good condition.
- Use the browser functions to print documents, and save Web pages to disk at every stage of each transaction.

48.6.10 Online Gambling

- Do not gamble with money you cannot afford to lose.
- Do not gamble online, except at well-known sites.
- If you do gamble online, do not gamble with money at sites hosted outside your own country.
- Do not give your credit card number to online gambling centers that are outside your own country.

- Before you gamble online, do some research to find out if there have been complaints about that casino. Contact your Better Business Bureau, or equivalent, and see if you can find friends or acquaintances who have played on the site you are considering.

48.6.11 Preventing Malware Infections

- Keep your virus strings up to date (automatic daily updates are good).
- Do not download or use software that purports to help you break the law or cheat people and businesses.
- Do not download or use software that has been copied without permission or in violation of license restrictions. That is software piracy, copyright infringement, or plain theft.
- Do not execute software that anyone sends you through email even if you know and like the person who sent it to you. Just because the person is nice does not mean he or she is qualified to inspect programs for safety.
- Before sending someone an attachment such as a picture or any other kind of file by email, let your recipient know what to expect via a preliminary message; if you do not know the person personally, send an email requesting permission to send the attachment.
- Never open attachments you have received without advance notice, regardless of who sent them or what the subject line or text says. Be especially suspicious of generic subjects such as "FYI" without details or "You'll like this." If you are really curious about the attachment, phone or email the supposed sender to find out whether it is legitimate. However, remember that you should not run programs you receive as attachments, regardless of what the sender thinks.
- Do not forward programs, even reliable programs, to anyone; instead, tell your friends where to download useful programs from a trustworthy source, such as a legitimate Website.
- Before sending anyone a Microsoft Word document as an attachment, save the document as an RTF file instead of as the usual DOC file. RTF files do not include document macros and therefore cannot carry macroviruses.
- Disable macros in Microsoft Word.
- Use the options offered by your email client to shut off automatic opening or execution of attachments.
- Do not circulate virus warnings; if you insist on doing so, personally check their validity on any of a number of virus-information and hoax sites on the Web.

48.6.12 Guarding against Spyware

- Before installing freeware or adware, read the terms and conditions carefully to see if they currently include language permitting automatic transfer of information to the supplier or to third parties. Be aware that these contracts often include language authorizing the supplier to change the terms and conditions at any time and without notifying you.

- Install and use a spyware scanner and removal program, such as the free Ad-Aware program from Lavasoft, PestPatrol from Computer Associates, or ZoneAlarm firewall.
- If you are particularly irritated by spyware, install a real-time spyware monitor and blocker such as those just mentioned.
- Support legislative attempts to force software manufacturers to disclose their use of spyware.

48.6.13 Junk Email

- Do not buy products or services from anyone who has sent you junk email. If the company is unprofessional or inconsiderate enough to use such methods of advertising, it does not deserve either your business or your trust.
- Do not assume that the FROM address is correct, because often it is either nonexistent or, worse, fraudulently misrepresents the origin by pointing to a legitimate business that is completely innocent of wrongdoing. Never bombard the owner of a FROM address with multiple copies, or even one copy, of abusive email. Such messages, known as *mail-bombs,* will probably reach the wrong target—some innocent addressee.
- Never respond to the address listed for removal from an email distribution list unless you initiated the contact or are confident that you know the organization that sent you the message (e.g., publications you already subscribe to). Since bounces (returned email due to bad addresses) never reach them and there is no incremental cost for sending out addresses to unwilling people, these operators really do not care how you feel about the junk they send. Therefore, the unethical people who send junk email use the REMOVE function primarily to harvest correct email addresses so they can sell them to someone else.
- Even if you trust the organization that sent you a junk email, never click on a link contained in the message. Instead, visit the company's Website and request removal from their official contact address, which any reputable company has.
- Do not visit the URLs listed in junk email messages. Some of them are deliberately mislabeled and may bring you to offensive Websites.
- If you *really* feel angry about a particular email and it has a dropbox (a real address in the body of the message where you are supposed to reply), then if you have nothing better to do, you may want to send a copy of the spam to the appropriate address (usually in the form abuse@ISPname.domain where you have to fill in the variables *ISPname* and *domain*) address running the dropbox. However, the chances are high that your message will be one of hundreds or thousands of similar reports.
- Do not send any junk email yourself. Encourage those around you (friends, neighbors, children) not to send junk email either.

48.6.14 Mail Storms. Here are some simple suggestions for reducing the likelihood of mail storms:

- Minimize the use of automated responses on your email accounts.
- If you do autoforward your email, do not let your target mailbox fill up.

- If you are receiving autoforwarded email from your primary mailbox, do not autoforward back to the original mailbox.
- Email system administrators should receive exception reports identifying accounts with excessive numbers of email messages or excessive traffic, so that they can investigate for mail storms.
- Firewalls that inspect the content of email messages should be able to react to an excessive number of bounce messages from a single originating address by deleting the traffic or informing the system administrator of a likely mail storm.
- Managers of unmoderated lists should configure a FROM address different from the address that participants use to post messages to the list.
- Users of list servers who want to send personal messages should reply to the sender, not to the entire list.

48.6.15 Detecting Hoaxes. Key indicators that a message is a hoax:

- Use of exclamation marks. No official warning uses them.
- Use of lots of uppercase text, typical of youngsters.
- Misspellings and bad grammar.
- No date of origination or expiration.
- Inclusion of words like "yesterday" when there is no date on the message.
- References to official-sounding sources such as Microsoft, Computer Incident Advisory Capability (CIAC), Computer Emergency Response Team Coordination Center CERT-CC) but no specific document URLs for details. URLs for a site's home page do not count.
- No valid digital signature from a known security organization.
- Requests to circulate widely. No such request is ever made in official documents.
- Claims that someone is counting the number of email messages containing copies of the hoax.
- Threats about dire consequences if someone breaks the chain by refusing to forward the message.
- Claims of monetary rewards that make no sense. For example, the Disney organization will send you $5,000—for forwarding an email message.
- Use of complicated technical language such as "*n*-th dimensional infinite complexity control loops" that do not make sense.
- Claims of damage to computer hardware from viruses or other computer software.

48.6.16 Get-Rich-Quick Schemes

- Remind everyone to use common sense: Earning lots of money with little or no effort usually results in uncovering something impossible or illegal.
- Teach users the mantra of the skeptic: "If it sounds too good to be true, it usually is."
- Explain how dangerous it is to get involved with criminal schemes like using stolen or falsified credit cards. Talk about the victims of such fraud: everyone who

pays higher interest rates on unpaid credit card bills and innocent shopkeepers who lose merchandise to e-commerce crooks.

- Especially when talking to children, discuss Internet-mediated theft in the same terms as you discuss shoplifting. Explain how commerce works; point out that everyone suffers from all kinds of theft, including electronic shoplifting.

48.6.17 Hacking

- Contact your local FBI office and find out if they can send a speaker to your company or to a local meeting of a professional security association for a discussion of computer crime.
- If you or specific authorized staff (e.g., from the security group) do visit Websites that support criminal hacking, be sure to use a personal firewall and set the parameters to deny access to personal information and to refuse cookies and active code (ActiveX, Java) from such sites.

48.7 CONCLUDING REMARKS. This chapter focuses specifically on the use and abuse of Internet and email resources. However, it becomes clear that both of these technologies are simply extensions of the human being behind the computer. Whether it is pornography, online gambling, deceitful emails, or simply posting inappropriate material to a public Website, the potential for damage to an individual, family, or organization is high. Taking a proactive stance through education and awareness is one major tool to combat these deceptive and unethical practices.

Employers have an ethical, and in many cases legal, responsibility to develop and implement policies around the appropriate use of the Internet and email at work. Unfortunately, simply putting the policies out for employee consumption is not enough. Employers must continually remind employees about both the dangers of misuse and the potential consequences to their employment. And when an employee chooses to violate the policy, employers must have a clearly defined process for encouraging proper behavior.

The scope of impact for these issues does not simply end when the employees leaves the office. Because of the widespread use of Internet and email into nearly every facet of our lives, taking the message home to the family is an important responsibility for everyone. Child predators use the Internet to prey on unsuspecting or naïve children, in an effort to exploit them for whatever immoral activity they wish. Parents then have an obligation to put their own family policies in place regarding what is, or is not, acceptable use of the Internet and email in the home.

Unfortunately, there is no easy answer to the problems described in this chapter. Both the Internet and email are neither good nor bad. They only become good or bad by the users and their actions. As technology continues to increase in speed, and in the ability to store more data in less space, everyone must take an active role in protecting each other at the corporate level and in the home.

48.8 FURTHER READING

Blanpain, R., and M. Van Gestel. *Use and Monitoring of Email, Intranet, and Internet Facilities at Work: Law and Practice*. The Hague: Kluwer Law International, 2004.

Cavanaugh, C. *Managing Your E-Mail: Thinking Outside the Inbox*. Hoboken, NJ: John Wiley & Sons, 2003.

Criddle, L. *Look Both Ways: Help Protect Your Family on the Internet*. Redmond, WA: Microsoft Press, 2006.

Fraser, J. M. *Email Etiquette for Business Success: Use Emotional Intelligence to Communicate Effectively in the Business World*. [self-published], 2011.

Flynn, N., and R. Kahn. *E-Mail Rules: A Business Guide to Managing Policies, Security, and Legal Issues for E-Mail and Digital Communication*. New York: AMACOM, 2003.

Goldsmith, J., and T. Wu. *Who Controls the Internet? Illusions of a Borderless World*. New York: Oxford University Press, 2006.

Holtz, S. *Corporate Communications: A Guide to Crafting Effective and Appropriate Internal Communication*. New York: AMACOM, 2004.

Jovin, E. *Email Etiquette for Business Professionals*. New York: Syntaxis Press, 2007.

Payne, T. P., and A. Proops. *Employee Surveillance*. Jordan Publishing, 2013.

Spinello, R. A. *Regulating Cyberspace: The Policies and Technologies of Control*. Westport, CT: Quorum Books, 2002.

Stanton, J. M., and K. R. Stam. *The Visible Employee: Using Workplace Monitoring and Surveillance to Protect Information Assets—Without Compromising Employee Privacy or Trust.* Information Today, 2006.

Willard, N. E. *Cyberbullying and Cyberthreats: Responding to the Challenge of Online Social Aggression, Threats, and Distress*, 2nd ed. Champaign, IL: Research Press, 2007.

48.9 NOTES

1. Parts of this article are based on materials originally published by M. E. Kabay in a series of articles in his Network World Security Strategies column between 2001 and 2011 and in older writings. To avoid cluttering the text with nonessential endnotes, no quotation marks or references are provided for such material. A complete archive is available. www.mekabay.com/nwss
2. Gaudin, "Insider Threats Giving IT Execs Nightmares," eSecurityplanet, November 4, 2005, www.esecurityplanet.com/prevention/article.php/3561761
3. http://search.dhs.gov/search?query=daily+report&affiliate=dhs
4. J. Leyden, "The Enemy Within: Geeks, Squatters and Saboteurs Threaten Corporate Security," *The Register*, December 15, 2005, www.theregister.co.uk/2005/12/15/mcafee_internal_security_survey
5. S. Gaudin, "Sarah Palin fans try to rewrite history on Wikipedia: Former governor's supporters battle with Wikipedia editors over Paul Revere page," *Computerworld*, June 6, 2011, www.computerworld.com/s/article/9217359/Sarah_Palin_fans_try_to_rewrite_history_on_Wikipedia
6. *Parker v. C. N. Enterprises Order, District Court of Travis County, Texas*. Final Judgment, www.loundy.com/CASES/Parker_v_CN_Enterprises.html
7. Information on the Ponzi scheme is taken from James Trager, *The People's Chronology* (Henry Holt & Co. 1995, 1996), made available through the Microsoft Encarta 2007 CD. All rights reserved.
8. C. Tran, "Reading Is Believing." ScienceNow, July 19, 2005; http://news.sciencemag.org/sciencenow/2005/07/19-01.html
9. A. Hamilton, "Arizona Lottery Pick 3 Random Number Bug," *RISKS Forum Digest* 19, No. 83 (1998); http://catless.ncl.ac.uk/Risks/19.83.html=subj5#subj5

10. E. Werner, "Feds Say Vague Law Makes Ban on Internet Gambling Tough to Enforce," Associated Press, April 2, 2008; www.sddt.com/News/article.cfm?SourceCode=20080402fac
11. K. Young, "Internet Addiction Test," Center for Internet Addiction Recovery, 2008, www.globaladdiction.org/dldocs/GLOBALADDICTION-Scales-InternetAddictionTest.pdf
12. P. Bates and M. Fain, "Cheating 101: Paper Mills and You," Coastal Carolina University, 2000, www.coastal.edu/library/presentations/papermil.html
13. J. Rush, "Downing St Dossier Plagiarised," Channel Four News, February 6, 2003, www.channel4.com/news/articles/politics/domestic_politics/downing%2Bst%2Bdossier%2Bplagiarised%2B%2B%2B/253293.html
14. A. Larson, "Sexual Harassment Law," ExpertLaw, 2003, www.expertlaw.com/library/employment/sexual_harassment.html#FN1
15. M. Potok, "Internet Hate and the Law: The First Amendment Protects Most Hateful Speech on the World Wide Web, but There Are Exceptions," Southern Poverty Law Center Intelligence Report 97 (Winter 2000), www.splcenter.org/get-informed/intelligence-report/browse-all-issues/2000/winter/internet-hate-and-the-law

CHAPTER 49

IMPLEMENTING A SECURITY-AWARENESS PROGRAM

K. Rudolph

49.1 INTRODUCTION 49·2

49.2 KEY CONCEPTS 49·3

49.2.1 Learning Continuum 49·3
49.2.2 Awareness: Common Sense or Common Knowledge? 49·5
49.2.3 Focus on Behavior 49·6

49.3 CRITICAL SUCCESS FACTORS 49·8

49.3.1 Information Security Policy 49·9
49.3.2 Senior Level Management Commitment 49·10
49.3.3 Resources with Security, Communications, and Training Expertise 49·11
49.3.4 Visibility, Audience Appeal, and Participation 49·12
49.3.5 Destination and Road Maps 49·13
49.3.6 Common Challenges and Mistakes 49·15

49.4 TOPICS 49·17

49.5 TECHNIQUES FOR GETTING AND HOLDING ATTENTION 49·18

49.5.1 Images 49·19
49.5.2 Video 49·19
49.5.3 Surprise, Novelty, and Expectation Failure 49·21
49.5.4 Conversational Style 49·22
49.5.5 Analogies and Examples 49·22
49.5.6 Stories and Empathy 49·23
49.5.7 Currency 49·23
49.5.8 Credibility 49·24
49.5.9 Social Proof 49·25
49.5.10 Accessibility, Diversity, and Culture 49·26
49.5.11 Spaced Repetition 49·27
49.5.12 Pretest and Refine Messages and Methods Before Distributing Them 49·27

49.6 TOOLS 49·29

49.6.1 Intranet Website 49·29
49.6.2 Social Media and Crowd Sourcing 49·30
49.6.3 Videos and Podcasts 49·30
49.6.4 Compliance Statements 49·30
49.6.5 Sign-on Messages, Networked Screen Savers 49·31
49.6.6 Publications 49·31
49.6.7 Posters and Digital Signage 49·31

49.6.8 eLearning Courses 49·31
49.6.9 Classroom Training and Clickers 49·32
49.6.10 People Penetration Tests and Spear Phishing Exercises 49·33
49.6.11 Contests and Incentive Prizes 49·34
49.6.12 Awards and Recognition 49·36
49.6.13 Human Libraries 49·37
49.6.14 Volunteer Activities 49·38
49.6.15 Inspections and Audits 49·38

49.7 EVALUATION AND METRICS 49·38
49.7.1 Baseline 49·39
49.7.2 Metrics 49·39

49.8 CONCLUDING REMARKS 49·41

49.9 GLOSSARY 49·42

49.10 NOTES 49·43

49.1 INTRODUCTION. An active security-awareness program is not a luxury: It is a business necessity. Losses from security failures are growing at an increasing rate. Some examples of such failures include:

- A 2012 New Year's Eve California office building burglary led to the March collapse and bankruptcy of a national medical records firm because it allowed medical records and social security numbers of14,000 people to be exposed[1];
- As of March 2012, BlueCross BlueShield had spent $18.5 million resolving a 2009 hard drive heist, not including the value of the data itself[2]; and
- An attack on Sony's PlayStation Network exposed personal details of 90,000 customers, which analysts say will cost as much as $2 billion to fix.[3]

According to the Government Accountability Office, the number of reported security breaches increased from 5,503 in 2006 to 41,776 in 2010, an increase of 650 percent.[4] The number of compromised records reported in 2011 was 174 million in Verizon's 2012 data-breach report.[5]

Information security is the translation of the instinctive ability to recognize and react to physical threats into an ability to recognize and respond to threats to information assets (e.g., hardware, software, data, and information).

The purpose of an information-security awareness program is to prevent loss and to ensure compliance with laws and regulations. Security awareness helps to instantiate security policies—to convert theoretical advice and instructions on protecting information assets into observable, positive behavior.

Information-security awareness leads people to:

- Pay attention to what is happening around them;
- Recognize suspicious circumstances that may be security violations;
- Know the initial actions to take in response to their suspicions; and
- Take the appropriate actions in response (people often know what they should do, but are reluctant to get involved).

Security awareness is the result of activities, tools, and techniques that help a target audience focus on identifying what in their environment has value that must be protected (physical assets such as a laptop or mobile phone and intangibles such as data/information) and what they can do to provide that protection. For a business, the target audience must include owners, employees, contractors, suppliers, partners,

customers, and any other individuals who have or require access to the organization's information or information systems. A well-trained workforce is inarguably the most cost-effective security control.

The most important security-awareness messages are:

- What must be protected and why should I care?
- Why am I important to security?
- What do security incidents look like?
- What do I do about security?

Effective awareness programs motivate people, provide measurable benefits, and measure those benefits. Awareness materials must compete to gain people's attention. They must also be tailored to an audience, their work environment, and the technologies they use to achieve maximum impact. This chapter outlines a practical approach for implementing an effective security-awareness program.

49.2 KEY CONCEPTS. Key concepts for this chapter are that security awareness is part of a learning continuum, awareness relates more to common knowledge than common sense, and effective security awareness focuses on behavior modification.

49.2.1 Learning Continuum. "Information Security Training Requirements: A Role- and Performance-Based Model (Draft)," National Institute of Standards and Technology (NIST) Special Publication (SP) 800-16 Revision 1, addresses security awareness and role-based training. This document defines an information-security learning continuum:

- Awareness,
- Awareness training,
- Role-based training, and
- Education and professional development.[6]

Awareness applies to all employees whether or not they have access to information systems (e.g., a groundskeeper who notices someone carrying boxes out of the office after hours needs to recognize that he or she may be witnessing a security incident and must know how to report the potential incident). Awareness training addresses security basics and literacy, and serves as a transition from awareness to role-based training.

SP 800-16 REV1 states:

> Awareness training strives to build in an organization's information system user population a foundation of information-security terms and concepts upon which later role-based training, if required, can be based. Awareness training informs users of the threats and vulnerabilities that impact their organization and personal work environments by explaining the "what" but not the "how" of security, and communicating what is and what is not allowed. Awareness training not only communicates information-security policies and procedures that need to be followed, but also provides the foundation for any sanctions and disciplinary actions imposed for noncompliance. Awareness training is used to explain the rules of behavior for using a department's or agency's information systems and information and establishes a level of expectation on the acceptable use of the information and information systems.[7]

Awareness programs, as used in this chapter, include the first two levels of the continuum: awareness and basics and literacy. Awareness intends to focus attention on security. Awareness sets the stage for additional learning by changing individual perceptions and organizational culture so that security is recognized as critical and necessary. Security failures can keep individuals from successfully completing their work and can threaten organizational survival. Security-awareness activities have the following characteristics:

- Learning tends to be short term, immediate, and specific.
- Learners are information recipients.
- Learning can occur at the same time everywhere throughout an organization, and it can be continuous.
- Awareness activities are directed at broad audiences with attractive, attention-getting techniques, similar to those used in advertising and cause-marketing.

Basics and literacy applies to workers who use an organization's computers or access nonpublic information. Basics and literacy should direct the workforce regarding compliance with security controls and appropriate responses to attacks in process. Characteristics of basics and literacy are:

- Basics and literacy activities are more formal than awareness activities. The purpose of basics and literacy is to build knowledge and to change attitudes. Specifically, these activities "promote personal responsibility and positive behavioral change throughout an organization's information and information system user population, beyond what is disseminated in the organization's basic awareness efforts."[8]
- Basics and literacy activities use a variety of techniques to address different learning personalities and styles.
- Basics and literacy activities often include a course or presentation and a statement of acceptance of responsibilities.
- Basics and literacy activities typically start with an employee orientation, usually completed before the employee is placed in a work environment.
- Basics and literacy activities include periodic refresher activities on at least an annual schedule.

Awareness materials are generally broad in coverage, but limited in depth (that is, awareness covers a lot of ground, but does not dig very deep holes). Role-based training and education and professional development apply to staff with significant security-related roles or functions, including:

- Executives, such as the chief information officer (CIO) and the chief information security officer (CISO)
- Information system security officers (ISSOs) and staff
- Program and functional managers (e.g., system owners, information owners, network administrators, system administrators, security administrators)
- Application designers and developers

Role-based training takes longer than awareness, and involves establishing and enhancing skills and competency for those involved in functional specialties (e.g., management, systems design, and acquisition). Training is provided selectively based on an individual's job functions (roles) and is most effective when tailored to the business environment. Education and professional development are appropriate for those pursuing a security career, and include college-level courses and professional certifications, such as the Certified Information Systems Security Professional (CISSP) and the Certified Data Protection (CDP) specialist. Role-based training and education and professional development are beyond the scope of this chapter.

49.2.2 Awareness: Common Sense or Common Knowledge? Ira Winkler, president of the Internet Security Advisors Group, writes, "The fundamental issue is that of common sense vs. common knowledge. You cannot expect people to behave with common sense if they do not have a common knowledge."[9] Consider the following security-awareness messages. Are the messages below common sense?

- Do not share your password
- Do not discuss sensitive or protected information in public
- Create long, strong passwords
- Report actual or suspected incidents
- Delete email chain messages

Some would say that these messages are nothing but common sense. If these security messages are common sense, then why do we need to publish tips telling people not to share their passwords, not to let others tailgate into secure work areas, and not to talk about sensitive information in public? Perhaps these messages are "common sense waiting to happen." Or, possibly, these are common-sense items only within a particular environment or background, such as individuals who are computer literate.

In 2007, IRS workers familiar with the policy of not sharing passwords, disclosed their passwords because they did not understand that changing their password to one provided by a caller was the same as disclosing it.[10] An IRS audit group, posing as computer help desk staff, called 102 IRS employees asking each to help them "correct a computer problem" by providing their user name and temporarily changing their password to one the caller suggested.

- Sixty-one of the 102 employees did as requested.
- Managers were more lax than nonmanagers.
- A follow-up survey asking why the employees gave up their passwords so easily found that about one-third believed what they had been told by the unknown caller.
- *Ten percent believed that changing their password was not the same as disclosing it, which they knew was against the rules.*[11]

The 10 percent of respondents who did not recognize their actions as disclosure demonstrate a need for improved security-awareness activities. While most (but unfortunately not all) security professionals know to never change their password to one that someone else knows, this principle is not common knowledge among all computer

users. "When there are widespread problems, there is clearly a failure in how the security community is delivering the message."[12] An awareness program should instill common knowledge, and by using material that grabs attention and maintains interest, transform that common knowledge into common sense.

49.2.3 Focus on Behavior. An organization's workforce is generally among the first to be affected by a security incident. Their compliance with security policy can make or break a security program. A staff that is security-aware can detect and prevent many incidents and mitigate damage when incidents do occur. Thus, the need for awareness programs that focus on behaviors.

The need to focus on behaviors is critical, especially with younger workers. "Seven out of ten young employees who are aware of their companies' IT policies acknowledge breaking those rules with varying regularity, according to a Cisco survey of more than 2,800 college students and young professionals in 14 countries."[13]

The most common reasons for this are:

- The employees' belief that they aren't doing anything wrong (33 percent)
- The need to access unauthorized applications for their jobs (22 percent)
- Lack of enforcement (19 percent)
- Lack of time to think about policies (18 percent)
- Inconvenience of adhering to policies (16 percent)
- Forgetting to follow policies (15 percent)

Nearly two-thirds (61 percent) said that the responsibility for protecting information and devices is on IT or service providers and not on individual employees. As this example illustrates, security is too important to sacrifice to the status quo.

Security-related behaviors can be classed as good, bad, or ugly[14]: Good behavior complies with the letter of the law or better, the spirit of the law; bad behavior includes naïve mistakes or dangerous tinkering for example, sharing a password, writing a packet-spoofing application to test one's programming ability, or scanning the organization network to see how it works; and ugly behavior consists of intentional misuse or destruction for example, building script that disables other users' terminal sessions, forging email header information, using a file decryption program to access trade secrets without authorization, or introducing a Trojan horse program into the network.

Within an organization, good security behaviors include compliance with security policy, such as:

- Releasing nonpublic information only with appropriate authorization
- Promptly reporting a potential security vulnerability such as a lost mobile device
- Politely terminating and then reporting a suspected social engineering attempt
- Creating and using strong, unique passwords for critical systems

Bad security typically experienced includes:

- Sharing passwords
- Deploying a wireless network gateway that allows noncompany personnel to use the company's network

- Setting up a packet-spoofing application to test the user's programming ability
- Setting up a network monitoring scanner on the user's PC

Ugly security behaviors include intentional misuse such as:

- Building a script to disable other users' terminal sessions
- Forging email header information
- Intentionally introducing malicious code into the organization computing infrastructure
- Using someone else's email to send messages

Behaviors directly affect the ability of the organization to meet its business objectives. Good behaviors promote business objectives and allow resources to be focused on achieving organizational goals (e.g., improved profitability and reduced costs). Bad and ugly behaviors result in wasted resources and loss of workforce focus on business objectives (that is, increased costs and reduced performance). Bad and ugly behaviors, carried to their extreme, can result in organizational failure (e.g., intentionally compromising customer financial records could result in a financial liability so large that the company goes bankrupt). Exhibits 49.1 and 49.2 illustrate a table showing measurable end user security behaviors, whether the behaviors are good, bad, or ugly, and how they could be measured.

Effective security-awareness programs encourage people to treat mistakes as "portals of discovery" where they can learn how to avoid similar mistakes. Employees should recognize that it is in everyone's best interest to limit damage from a mistake and, more importantly, to learn from it, and report it quickly rather than fixing it quietly. Mistakes often have side effects, and pretending that a mistake didn't happen is dangerous. In a *Harvard Business Review* interview, former Toyota chairman Katsuaki Watanabe said, "Hidden problems are the ones that become serious threats eventually. If problems are revealed for everyone to see, I will feel reassured. Because once problems have been visualized, even if our people didn't notice them earlier, they will rack their brains to find solutions to them."[15] Sharing what you learned can prevent loss. In her research on learning in hospitals, Amy Edmondson of Harvard University discovered that the highest-performing nursing units had reported the largest number of mistakes. Not because they made more mistakes, but because they felt safe to report and share the ones they did make.[16] Learning the cause and how to avoid mistakes is vital to security.

Technologists typically try to control workforce behaviors by adding layers of technical controls. As demonstrated by the increasing number of reported information compromises, technological approaches alone cannot solve a people problem. The workforce needs access to data and computing functions to do their job. Technology cannot effectively distinguish between the instances when an individual employs a capability they are authorized to use to accommodate a good behavior versus a bad behavior. Technology cannot adequately address the human factor of intent.

The rise of mobile computing in the work environment has brought increased importance to the human factor. The network perimeter now extends from a defined physical area to wherever data might be at any given time (e.g., an employee's home, a laptop at an offsite meeting). Building good computing habits at home is as important, if not more important, than building those behaviors at work. Secure computing habits will

Security Awareness Program Metrics Internal User Behaviors *Examples to be tailored should be based on goals and then questions.*	G · B · U*	Survey	Observation	Helpdesk Incid. Rpts	Manual Tests or Audits	Software (Automated)
% of users recognizing a security event scenario	G	✓				
% of users susceptible to social engineering (compare mid-week at mid-day to Friday afternoon)	B				✓	
% of users revealing their password when tested	B				✓	
% of users activating a "test virus"	B				✓	
% of security incidents having human behavior as a major factor/involving behaviors covered in awareness materials	BU			✓		
Storage of unauthorized file content on desktop or network resources, such as audio, video, or other multimedia files	B					✓
Number of attempts to access inappropriate/blocked Websites	B					✓
Nonpublic information found in dumpsters outside of facilities	B		✓			
% of systems having unapproved software installed	BU				✓	✓
% of systems having unapproved hardware installed	B				✓	✓
% of emails (random sample) with inappropriate content	B				✓	
% of passwords visible or in common locations (e.g., under lamp)	B		✓			
% of PCs logged on and unattended	B		✓			
% of laptops, portable devices/media, sensitive data unsecured	B		✓			
% of laptops, portable devices/media, stolen (office/travel)	B			✓		
Number of attempts to use unauthorized resources, e.g., VPN	B					✓
% of emails sent via Internet containing nonpublic/sensitive data that are not encrypted	B				✓	
% of users wearing badges with picture facing out	G		✓			
% of monitors positioned to be easily seen from hallways, doors, or windows (especially on the ground floor)	B		✓			

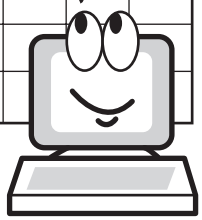


EXHIBIT 49.1 Awareness Metrics, continued as Exhibit 49.2

transfer across environments, as individuals recognize good behaviors are beneficial to them regardless of whether they are at home or at work.

49.3 CRITICAL SUCCESS FACTORS. Critical success factors for implementing a security-awareness program include:

- An information security policy
- Senior-level management commitment and buy-in, to demonstrate the importance of security
- Resources with subject matter, communications, and training expertise

Security Awareness Program Metrics Internal User Behaviors *Examples to be tailored should be based on goals and then questions.*	G · B · U*	Survey	Observation	Helpdesk Incid. Rpts	Manual Tests or Audits	Software (Automated)
% of users who challenge unknown visitor with no access badge	G		✓	✓		
% of users who open a test email with a questionable subject	B				✓	✓
% of users activating a "test virus"	B				✓	✓
% of users who click a link in a test email (instead of typing the URL into their browsers)	B				✓	✓
% of users responding to a test email via an "unsubscribe" link	B				✓	✓
% of crackable user passwords	B				✓	✓
% of user systems having spyware or malware installed	B				✓	✓
Number of incidents of unauthorized use of administrator privileges	U			✓	✓	
% of users sending Internet email to multiple recipients who do not use the BCC field	B				✓	
% who have actively acknowledged policies/security responsibilities	G					✓
Number of major findings from internal and external security audits	B				✓	
% viewing optional security materials in online courses	G					✓
% participating in contests, suggestion programs, bonus questions	G					✓

*__G·B·U =__ Good, Bad, or Ugly
G: Good behavior complies with the 'letter or the law' or better, the 'spirit of the law,' e.g., not releasing nonpublic information inappropriately, discovering and reporting a security vulnerability.
B: Bad behavior includes **naïve mistakes** or **dangerous tinkering —** e.g., sharing a password, deploying a wireless network gateway that allows noncompany personnel to use the company's network, setting up a packet spoofing application to test one's programming ability, or setting up a network monitoring scanner on one's PC.
U: Ugly behavior consists of **detrimental mususe** or **intentional destruction —** e.g., someone builds a special script that disabled other users' terminal sessions, forges email header information to make it look like someone else sent a message, uses a file decryption program to discover the contents of a file containing trade secrets, or intentionally introduces a Trojan horse program into the network.

(*Behavior categories inspired by "Analysis of End User Security Behaviors" - by Jeffrey M. Stanton, Kathryn R. Stam, Paul Mastrangelo, and Jeffrey Jolton, July 12, 2004.*)

EXHIBIT 49.2 Awareness Metrics, continued from Exhibit 49.1

- Visibility, audience appeal, and participation to address all subgroups within the workforce
- Destinations and road maps to guide and monitor program activities

49.3.1 Information Security Policy. Effective information security policies are in-place, credible, comprehensive, and current. Security objectives must be embodied in policies that clarify and document management's intentions and concerns. Policies are an organization's laws. They set expectations for employee performance and guide behaviors. Information security policies include statements of goals and responsibilities, and delineate what activities are allowed, what activities are not allowed, and what penalties may be imposed for failure to comply.

Effective information security policies show that management expects a focus on security. Well-defined security policies show what is expected of the workforce and make

it easier to take disciplinary action against those who ignore policy and compromise security.

A cohesive security-awareness policy provides credibility and visibility to the information-security program. It shows that management recognizes that security is important and that individuals should and will be held accountable for their actions. As Daryl White, chief information officer for the U.S. Department of the Interior, said, "You can't hold firewalls and intrusion detection systems accountable; you can only hold people accountable."[17] Credibility also requires that management back employees who do the right thing.

An awareness policy should address three basic concepts:

1. **Participation** in the awareness program is required for everyone, including senior management, part-time and full-time staff, new hires, contractors, and other outsiders who have access to the organization's information systems. New hires might be required to receive a security-awareness orientation briefing within a specific time (e.g., 30 days after hire) or before being allowed system access. Existing employees might be required to attend an awareness activity or take a course within one month of program initiation, and periodically thereafter (e.g., semiannually or annually).
2. Everyone will be given sufficient **time to participate** in awareness activities. In many organizations, security policy also requires that employees sign a statement indicating that they understand the material presented and will comply with security policies.
3. **Responsibility** for conducting awareness program activities is assigned. The program might be created and implemented by one or a combination of: the training department, and the security staff, or an outside organization, consultant, or security-awareness specialist.

49.3.2 Senior Level Management Commitment. Senior management must be committed to information security and visibly demonstrate that commitment by example (e.g., signing the awareness program or activity launch announcement, participating in awareness activities), providing an adequate budget, and supporting the security staff. Saying security is important but failing to follow organizational policies will have negative consequences.

Executives set the standard for organizational behavior. For example, in Colombia, when there was a water shortage, the mayor of Bogotá, Antanas Mockus, appeared on television programs taking a shower and turning off the water as he soaped, asking his fellow citizens to do the same. In just two months people were using 14 percent less water, a savings that increased when people realized how much money they were also saving because of economic incentives approved by Mockus. Water use is now 40 percent less than before the shortage.[18]

Organizational leaders must understand and support the program as well as provide oversight. Program responsibility has shifted in recent years from a collateral duty of a compliance or information-security officer to the highest levels of the organization.

Poor security measures can be costly in damage to the organization's brand or reputation, in impact on operations, and in actual and potential lawsuits. The media will not hesitate to report a security threat or breach. Such stories are a wake-up call and highlight the need for senior executive commitment to the security function.

Wise senior managers know that security is not just about reducing risk, it's also a tool to protect the organization's reputation. It also builds customer confidence and market valuations, and delivers a competitive edge. In today's e-commerce environment, effective information security can increase business and profits. Top management should understand that security is not solely a risk avoidance measure and should see the security-awareness program as a business enabler.[19]

Senior-management compliance provides credibility. If security policy prohibits installation of personal software on organization infrastructure, even senior executives must comply. Doing otherwise undermines the policy and creates a perception of inconsistence, unfairness, or unimportance. Senior managers must stand behind the organization's policies and the security staff charged with enforcing those policies. Consistent enforcement is especially important in areas where security and convenience conflict, such as changing passwords frequently or enforcing denial of system access for users who have not completed a required awareness refresher activity. In addition, human beings tend to imitate those with higher social status, so executives who see their superiors refusing to wear name badges will soon be doing the same—and the breach of security policy will propagate downward through the entire organization.[20]

Implementing an awareness program is always a management challenge. Senior managers generally appear to recognize the benefits of an awareness program, but are often still reluctant to allocate the financial and staff resources necessary to make it effective. Awareness programs must compete against other organizational needs. It is relatively easy to identify the cost of an awareness program, but it is difficult to quantify its benefits. Thus, awareness programs often lose when competing against programs where benefits are more tangible (e.g., programs that return a profit). The difficulty in quantifying benefits is a primary reason why the U.S. Government made maintenance of a computer security-awareness program mandatory for federal organizations[21] and why security awareness is required by laws and regulation in specific industries such as finance and healthcare.

Management resistance is often tied to viewing security awareness as "nice to have" but not as important as other needs vying for limited funds. Although the time and effort to build a strong security program is not trivial, it is far less than the time and effort required to deal with just one serious incident. Some security professionals recommend equating awareness with insurance policies. Insurance policies require continuous funding but are not often used (and the hope is not to use them); however, few organizations choose to forego those costs.

When common sense and fiduciary responsibilities are not enough, legal requirements provide another incentive for awareness programs. The Federal Information Security Management Act and the Computer Security Act require such programs for federal organizations. State and federal laws (e.g., the Sarbanes-Oxley Act and the Gramm-Leach-Bliley Act for financial information, and the Health Insurance Portability and Accountability Act for healthcare information) require security-awareness components in information-security programs for state agencies and public companies. Requirements to maintain effective security programs are also included in state and federal contracts that impact critical infrastructures (e.g., healthcare, public safety) and contracts that require retention of sensitive customer information.[22]

49.3.3 Resources with Security, Communications, and Training Expertise. Ideally, responsibility for an awareness program should be assigned to one individual who has a defined budget and has security subject matter expertise, communications skills (such as marketing), and knowledge of training principles. This

individual should be a member of senior management and performance of the awareness program should be a factor in his or her performance evaluation. Program activities and materials may be created and deployed by one or more individuals drawn from internal or external resources.

The mix of skills needed to successfully maintain a security-awareness program include:

1. **Subject matter expertise**—Individuals attempting to infuse security awareness into the business processes should be knowledgeable of those business processes. This ensures that awareness messages support the business and are delivered using terminology and content that are accurate and aligned with other organizational initiatives and activities (e.g., a professional security certification is valuable when working with a technical audience to establish credibility).
2. **Communications expertise**—Awareness activities often require translation of technical, security-oriented information into tips, discussions, presentations, posters, job aids, and other tools that average computer users can understand and apply in their daily work. Communications skills facilitate development of these awareness materials.
3. **Training expertise**—Training is its own discipline. People learn in different ways. Knowledge of training techniques, how people learn, and how to match training techniques to learning styles will improve the potential success of an awareness program.

The awareness team also needs to include an individual with deep information security experience. This individual serves to validate material for the general workforce and to help avoid the *curse of knowledge*. Once we know something—say, the melody of a song—we find it hard to imagine not knowing it. Our knowledge has metaphorically cursed us. We have difficulty sharing our own knowledge because we cannot readily relate to our audience's state of mind.[23]

The curse of knowledge was identified in a 1990 study by Elizabeth Newton, a graduate student at Stanford University. Study participants were divided into two groups: tappers and listeners. The experimenters chose 120 well-known songs such as "Happy Birthday" and the tappers tapped out the rhythm on a table while the listeners tried to guess the song. Before they started, the tappers were asked to predict listener success. Most predicted about 50 percent. The actual success ratio was 2.5 percent. The tappers conveyed the message successfully one time in 40, but thought that they would get the message across one time in two. This discrepancy resulted from the fact that the tappers hear the tune internally while they tap, but the listeners only hear what sounds like random taps.[24]

49.3.4 Visibility, Audience Appeal, and Participation. An effective awareness program cultivates a professional, positive, and visible image. A visible program demonstrates the value of the awareness activities, raises employee morale, and encourages the support of the general workforce. The more methods used to spread the message, the more visible the program. An awareness program that uses computer-based courses, videos, posters, acknowledgment statements, newsletters, contests, events, daily tips, and checklists will reach more people and have a greater impact than a program that consists of posting security policies to the organization's intranet and sending a memo advising staff to read the policies.

Everyone should receive sufficient time to participate in awareness activities; this includes those who are responsible for planning and developing awareness activities. These activities should occur on compensated time. Organizations that require employees to only obtain or develop security awareness and training on their own time effectively state that security is not important.

Security-awareness programs that show the organization's concern for employees' IT security well-being at home (for telecommuters and others who use computers at home) and while traveling are better received than programs that ignore such issues. Practical topics that cross between home and office might include "what to do if your Twitter account is hacked," "how to use Apple iOS5 securely," or "how to protect your personal email account." Whether the target audience is all end users or senior management, showing them how they will personally benefit from improved security awareness contributes to program success. Viewing security as a service, with the entire organization as a customer, highlights the importance of marketing security to management and staff.

Computer behaviors and habits from home transfer to work. Security professional Donna Mattick put it this way,

> Just knowing that my elderly relatives are using the Internet causes me to stay up at night worrying, but it also drives me to find ways to protect them automatically. On the other end of the scale, I have a teenage daughter who has an iPad, cell phone, and exposure to computers at school. She can find a way around every protection measure I put in place. So I have to stay current to keep up with her. … Children, parents, and elders all need to be cyberaware and we security professionals need to step up and help. For every person we educate how to stay safe we chip away at the criminals' ability to take advantage of us.[25]

49.3.5 Destination and Road Maps. When a psychologist was invited to give a talk at the Pentagon on managing time and resources, he decided to warm up the group of generals with a short exercise. He asked them all to write a summary of their strategic approach in no more than 25 words. The exercise stumped most of them. The only general who managed a response was one who had worked her way through the ranks and been wounded in combat in Iraq. Her approach was as follows: "First I make a list of priorities: one, two, three, and so on. Then I cross out everything from three down."[26]

Create a security-awareness program plan that contains these elements:

- A description of the organization and its IT culture (culture is the instinctive behavior of individuals within an organization), including assigned roles and responsibilities
- Program goals and the status of the organization's current efforts with a security baseline
- A determination of awareness needs by audience
- A description of methods and materials to be acquired, created, and/or modified
- A schedule showing actions to be completed and who is responsible for ensuring their completion (including program evaluation and updates)

49.3.5.1 Goals. A security-awareness program should have goals and a plan for achieving the goals that includes measureable criteria. The goals and objectives should be related to improvements in workforce performance and security

risk. For example, the following might be appropriate goals for a security-awareness program:

- Improve employees' ability to recognize potential threats and vulnerabilities
- Improve the level of compliance with company physical and computer security controls
- Reduce the occurrence of security failures resulting from employee action or inaction
- Reduce the severity of the security incidents that do occur

Specific, realistic, and measurable goals are best. If possible, establish a baseline prior to implementing your awareness program or launching a new campaign. A baseline defines where the organization stands with regard to its security-awareness efforts and program. It may be that the baseline is zero and a program has not been implemented. In that case, the organization may take a survey to find out how people in the organization view security and how familiar they are with security policies. The results of the survey would become the baseline. Where a program has been implemented, the organization may choose to document the level of awareness so that over time, other measurements can be taken to show changes.

Victor Basili of the University of Maryland developed an approach for metrics where the metric is created as the final step of a process called Goal-Question-Metric.[27] He recommends defining a goal, for example: "Goal—decrease inappropriate Website visits." Next, create a question that will indicate whether the goal is being met or not: "Are staff continuing to visit Websites that they should not?" Then, and only then, create a metric that will support the goal. The metric would be the number of attempts to access inappropriate Websites, such as illegal or pornographic material. This information can be extracted from Web filtering products. This approach offers several benefits:

- It leads to an automated metric
- The information for the metric is easy to collect
- The metric will give a constantly updated idea of what the organization's users are doing

49.3.5.2 Audiences and Messages. Awareness programs need to be planned to ensure they address the intended target audience in the appropriate manner. If you are directing your message to your total workforce, then you should consider consultants, contractors, subcontractors, vendors, suppliers, and other third parties. When your audience is employees, don't forget temporary hires and interns. If you divide your audience by technical skill, you can provide detailed information to your technical staff and less detail to your clerical staff.

Once you decide on your audience and how to segment that audience, you can develop messages and delivery methods specific to each segment. The following are audience breakdowns for a typical business:

- **Executives**—are generally interested in return on investment and risk reduction, have a moderate level of technical literacy, and are key targets for spear phishing and social engineering.

- **Information technology staff**—should be interested in building security into applications, networks, and systems, should have specific security responsibilities, and are expected to have moderate to high computer literacy.
- **Business users**—primarily focus on getting the business process completed regardless of the security impact. They can be easily frustrated when security controls are time consuming or inhibit "real work," but they can be motivated when they can see the benefits to themselves and their families, and will respond positively once they understand the value of security controls, especially if they handle sensitive or protected customer data.
- **New hires**—for those for whom this is their first "real" job, awareness will need to focus on the fundamentals (what needs protection, from whom or what, why they should care, and what is their role in the organization). For new hires who have been in the job market, awareness messages generally need to focus on how security is addressed in this organization.
- **Mobile device and smart phone users**—often associate risk with the smallness and low cost of the device instead of the adverse impact their compromise or loss could have on themselves as well as on the company.
- **People who travel**—whether they travel locally or abroad, travelers have removed themselves and their computer devices from the friendly surroundings of the office and home. An unattended device, in a hotel room, in the trunk of a car, or even at airport screening, is an invitation for theft and data compromise.

These messages should be delivered using a mix of tools (e.g., posters, screen savers, presentations, events, computer-based courseware, classroom training, and one-on-one training). The more times the same message is delivered using different approaches, the greater the likelihood the information will be retained.

49.3.5.3 Methods and Frequency.

Some companies use a perpetual calendar for their security program. The calendar is used to create a communications plan that covers the type and frequency of message by audience. A calendar-oriented plan shows what security-awareness materials and activities are produced monthly, quarterly, or annually. Calendars need to be periodically updated to add special events, such as security-awareness days or weeks.

A security-awareness program should be an ongoing effort. Some organizations offer a security-awareness orientation to new employees and regular reinforcement for all employees at various times throughout the year. Doing so provides spaced repetition of the material, and reinforces learning. Some organizations address security awareness on a monthly basis with newsletters, posters, screen savers, contests, surveys, and online modules. Other organizations offer awareness courses that are updated annually and provide reinforcement at various times, such as on November 30, International Computer Security Awareness Day. NIST SP 800-50, "Building an Information Technology Security Awareness and Training Program"[28] presents a detailed approach for establishing and maintaining a security-awareness program, including an appendix with a sample awareness program plan template.

49.3.6 Common Challenges and Mistakes.

Awareness program planning should consider common obstacles and constraints such as keeping management informed, changing material language to relevant business terminology, gaining union support, and overcoming audience resistance. The program should also be structured

to address significant, but less recognized challenges of diffusion of responsibility (where a person is less likely to act when others are present because they assume that others will take action or have already done so); and attenuation (when a message loses its strength and the learners tune out, usually because the message has been reused, overused, or doesn't capture the student's interest). Wildlife managers describe attenuation as "getting used to something we shouldn't." They say: "A fed bear is a dead bear."[29] When bears in parks get used to campers and carelessly stored food, they adapt, resulting in dangerous encounters or mailings and the destruction of the bear. Bad things can happen when your awareness messages lose their signal strength.

A common mistake is not fitting the program to the environment. Build an awareness program around your business environment. If you take materials from other organizations, be sure to tailor the content to your environment and target audience. Free training materials can have significant costs in terms of wasted audience time and can increase security risks by providing a false sense of security if the training doesn't meet your program objectives.

Security program planners often overlook the learners. This typically occurs because those charged with planning and building the program were selected for their technical competence, but have limited experience in selecting learning methods and techniques. Programs that don't consider the learner often fail. Learners should be able to relate to the awareness materials and apply the materials to their jobs.[30]

Joseph A. Grau, former chief of the Information Security Division at the Department of Defense Security Institute, believed in the importance of marketing security and often stated that customers actually pay for security services. For example, managers pay for enforcing the requirement to lock a classified document in a safe rather than leaving it on a desk, with labor hours. Other methods of payment are in the form of energy, attention, and concern for security matters, such as taking time to identify and report a potential security incident. Even egos are part of the payment for security. There is an "ego cost" when "scientists, researchers, technical specialists, engineers, and management personnel must refrain from communicating their successes to friends, family, and peers to protect sensitive, company private or classified information."[31]

Be sure that your awareness program plan clearly defines your objective (how to know when you are successful) and how you can monitor your progress toward your stated objective. Successful performance is always easier to achieve when you have a road map or plan that tells you where you are, where you are going, and how to tell that you have reached your destination. Abraham Lincoln related planning to sharpening an ax. If he was asked to cut down a tree in eight hours, he would spend the first six sharpening his ax. It takes great effort to chop down a tree with a dull ax, but cutting a tree down with a sharp ax goes much more quickly. Similarly it is much harder to create, manage, and measure the effect of an unplanned awareness program than a planned program with defined objectives, assigned responsibilities, and management direction.

Careful planning promotes awareness activities that elicit specific, positive responses. Flexible plans allow timely changes to address changes in the organizational structure, objectives, new technologies, and applicable threats and vulnerabilities. Flexibility also allows incorporation of relevant current events, events that use external sources to emphasize your security message. Microsoft's policy of issuing awards for help in capturing virus writers is evidence that security issues are now getting high-level attention.

49.4 TOPICS. In planning an awareness program, you should have an understanding of the topics that you want to address. NIST SP 800-16, "Information Security Training Requirements: A Role- and Performance-Based Model (Draft)," identifies the following security-awareness topics:

- Roles and responsibilities in information security
- Ways to protect shared data (e.g., encryption, backups)
- Examples of internal and external threats (e.g., social engineering, hackers)
- Malicious code (e.g., viruses, worms)
- Security controls
- Ways to recognize an information-security incident
- Principles of information security
- Passwords
- Social engineering
- Data backup and storage
- Computer viruses and worms
- Incident response
- Personal use and gain
- Privacy
- Personally identifiable information (PII)
- Identity theft
- Internet surfing
- Inventory control
- Physical security
- Spyware
- Phishing
- Scams and spam
- Mobile devices (e.g., laptops, smartphones, tablet computers)
- Portable storage devices (e.g., CDs, USB drives)
- Remote access
- Copyright infringement and software piracy
- Use and abuse of email
- Email do's and don'ts
- Peer-to-peer file sharing threats
- National security information systems, where applicable[32]

Although these topics may not be directly relevant to a specific organization, they do provide a catalog to which you can add or subtract based on your industry, technology, and assets that need protection. For example, healthcare providers should address personal health information (PHI), both paper and electronic (ePHI); organizations that accept credit cards should include topics linked to the Payment Card Industry Data

Security Standards (PCI DSS); and organizations that deal extensively with intellectual property (IP) should consider adding topics that include advanced persistent threats (APT) and related social engineering techniques.

Exploiting technology weaknesses to gain financial reward is a significant threat, but less obvious data thefts can cause substantial harm to your organization's reputation and require significant clean-up costs. Identity theft can affect organizations as well as individuals, and stealing the ideas for a new marketing program or product can damage a firm's ability to effectively compete for new business.

Including the impact of data theft is a key topic in your awareness list. The impact of data theft was often not well understood, even though it could result in a firm's bankruptcy. It wasn't too long ago that a court held that a data theft could not be prosecuted because nothing was actually taken. Today's judges are more knowledgeable and recognize that the knowledge the data conveys (e.g., a person's identity) can be more damaging than stealing a person's money. Linking your awareness messages to impact on the organization is an effective approach to explaining impact and increasing acceptance of your security message.

Another way to identify appropriate topics is to ask managers, helpdesk, and incident response personnel to identify recurring problems or review the problem reports to identify recurring issues. Awareness topics can then be structured to help resolve these problems by surfacing their cause and how to avoid them. Also, topics that are of personal relevance are good for gaining attention. Data mining, mobile device location awareness, cyberbullying, identity theft, travel precautions, and the latest frauds, scams, and malware are of interest to most computer users.

49.5 TECHNIQUES FOR GETTING AND HOLDING ATTENTION.

> *If you want to build a ship, don't drum up the men to gather wood, divide the work, and give orders. Instead, teach them to yearn for the vast and endless sea.*[33]
>
> —Antoine de Saint-Exupery

If you want to secure information and computer systems, don't dictate orders or make commandments for employees to follow. Teach them why they are crucial to the security process. Learning methods that are interactive, demonstrative, and rewarding get the most attention. Work with the brain to capture and maintain attention. Before you can teach your audience anything, you must have their attention. The strongest messages have a visual and visceral impact and use images, surprise, novelty, emotional involvement, and empathy.

> Most of the events that predict whether something learned will also be remembered occur in the first few seconds of learning. The more elaborately we encode a memory during its initial moments, the stronger it will be.[34]

Emotional context plays a large role in memory retention. Emotional arousal helps the brain learn. People remember things that they feel, such as empathy when they learn about a person who accidentally reformatted a hard drive and did not have a backup. (Are my files backed up? Could I accidentally reformat my disk?) Other feelings that awareness program materials and events can convey include surprise, curiosity, and satisfaction (e.g., correctly answering a difficult quiz question or solving a security puzzle). Scenarios are a good way to create empathy because they allow choice of action to be directly associated with consequences.

Anything that increases brain activity causes deeper learning. The more different types of brain activity involved, the better the results of the awareness program. In short, design and implement awareness activities and messages that use as much of the brain (especially both sides) as possible.

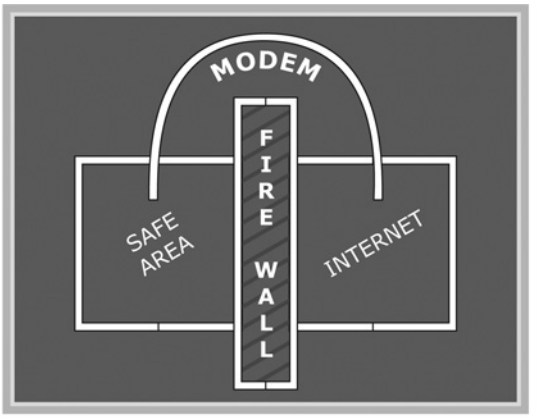
49.5.1 Images. "We do not see with our eyes. We see with our brains."[35] Half the brain is set up for visual processing. The more visual the input, the more likely it is to be remembered and recalled. This phenomenon is called the Pictorial Superiority Effect (PSE). The potency of the PSE was described definitively in 1976 by neuroscientist Douglas Nelson. "He and others have shown that our brains are essentially hard-wired for visuals—the very architecture of our visual cortex allows graphics a unique mainline into our consciousness."[36] People pay attention to color, size, and orientation. People pay special attention to objects that are moving. This is why animations (e.g., our computer game-oriented culture) are so effective.

Imagery stimulates both verbal and visual representations. Language is primarily processed through only the verbal channel. Experiments have shown that imagery activates multiple, powerful neural pathways of memory recall. While our access to raw information has grown, our time to process this information has declined (we are reduced to communicating in sound bites and 140 character Tweets). This places a premium on meaning-making. Given our brain's preference for the visual and the current complexity of our world, "we've learned that the very best shortcuts usually come in graphical form … consequently, today's visual storytellers have considerable power."[37]

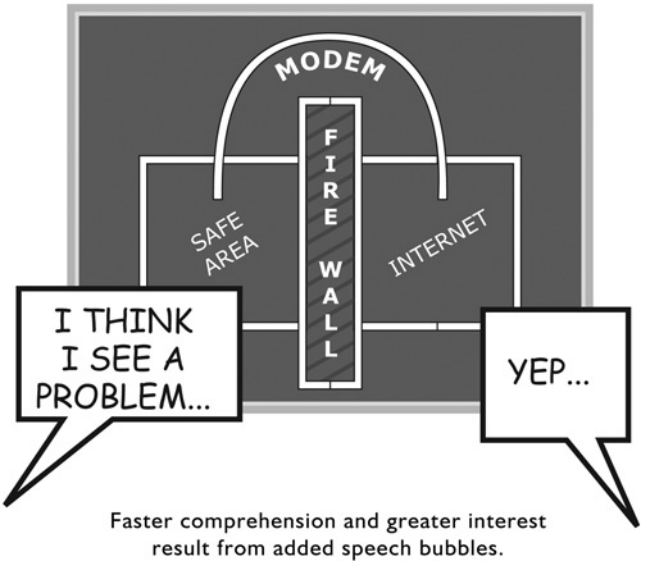
Graphics that incorporate the message in the image (e.g., speech boxes) are more effective than graphics that described the message or image in an accompanying narrative or that have the message as a caption or title below the image. Exhibit 49.3 shows two images of a vulnerability caused by a modem. Learners were observed to glance at the first one for a short period of time and many "tuned it out." When the speech boxes were added, learners looked at the image longer and expressed excitement, saying, "Oh, I get it. The modem goes around the firewall and circumvents the control."

49.5.2 Video. While some educational professionals say the future of online learning is gamification[38] (learning through games), video is arguably the future of online security awareness. YouTube is the second most popular search engine online—today. Humanity watches more than 80 million hours of YouTube every day, according to Chris Anderson, TED[39] founder. Our brains are wired for video more that print. Video offers a greater density of information. Anderson points out that print and reading are relatively new compared to face to face communication. Print was scalable, which was an advantage that has been overtaken by video. Anderson's TED talk addressed crowd-accelerated innovation and learning where cycles of improvement are driven by people watching Web video. He uses the examples of street dancers and TED talks. As a result of being able to see what others in their fields are doing, people are stepping up their game. Anderson identified three concepts that fuel accelerated learning and performance:

1. A crowd, such as a global Internet community—the bigger the crowd, the more potential innovators there are, as well as commenters, trendspotters, cheerleaders, skeptics, mavericks, and super-spreaders. These people are creating the ecosystem through which innovation emerges.



Graphic showing vulnerability. "Tuned out" by many learners.



Faster comprehension and greater interest result from added speech bubbles.

EXHIBIT 49.3

2. Light, clear, open visibility of what the best are doing—this empowers others to participate. A light shines on the innovators, either directly through comments, ratings, email, Facebook, and Twitter, or indirectly through numbers of views and links (that point Google there).
3. Desire for social status—where the best walks tall and is recognized—"You might just be a kid with a Webcam, but if you can do something that goes viral, you can be seen by the equivalent of sports stadiums crammed with people." This global recognition drives huge amounts of effort. The light and desire are self-fueling and attract new people to the crowd.[40]

Here's an example of how well video can work for security awareness described by Chip and Dan Heath. Russ Berland was tasked with redesigning BearingPoint's ethics and compliance training program. He inherited a code of conduct which might have

been repurposed from a law firm. A principal challenge was the need to influence the behavior of employees across the country, operating in different organizational cultures. Berland interviewed associates about real-world "gray areas" and uncovered dramatic stories of strained relationships and ethical quandaries. This inspired his team to create a humorous fictional series based on *The Office* featuring a fictional company designed to be the "evil doppelganger" of BearingPoint. The fictional company, Aggrieva, used the motto: "Aggrieva says yes when everyone else says no." Berland hired a film maker and shot 10 short episodes over a weekend. The films included topics such as bosses "hitting on" subordinates, teams misrepresenting their expertise, and managers trying to pass along inappropriate expenses to the client.

The episodes created a sensation and employees said that this was the best training they had ever had. The characters and situations became part of the company's vocabulary. "New episodes debuted each Monday, but employees were so ravenous for the next episode that they started tracking them down on the company's staging server, where the videos were posted on the preceding Friday. Thousands of employees watched the videos before they were released."[41]

The videos started conversations about ethics and compliance, and after the videos aired more people called the hotline to discuss difficult topics and situations. The videos were so well-received that people chose to watch them. Awareness programs should aim to create or locate materials so engaging that people are eager to watch.

49.5.3 Surprise, Novelty, and Expectation Failure. Vital information about potential threats and resources is likelier to be identified from things that are new or unfamiliar. Nature ensures that all living creatures react to novelty and change because novelty and change often results in danger. A swerving car on the highway, a jump in your bad cholesterol, or a drop in a stock's value rivets your attention and jangles your nerves, events which prime you to protect yourself from harm. Basically, our brains are surprise detectors.[42]

Use surprise and expectation failure to deepen security-awareness experiences. Humans crave novelty from evolution and for survival and will respond to the unique and unusual. Anything that is counter-expectational will tweak the arousal–adaptation cycle.

Ask questions such as, "Did You Know?" Ask learners with mobile phones how long their data (multimedia pieces of communication like photos, videos, and texts) is stored by their cell phone network provider and if their data can be sold to third parties. Not many people know that nearly four years' worth of their digital identity is stored by AT&T, which holds the data for 84 months. Verizon holds this data for 12 months, Sprint for 24 months, and T-Mobile for 60 months. "28,000 MMS messages are sent into the world every second, and cell phone companies record much of the metadata that travels with them, like location, receiver identity, amount of data transferred, and the cost of the transmission. The average user has 736 pieces of this personal data collected every day."[43]

Another innovative idea from the mayor of Bogotá, Antanas Mockus, was to use mimes to improve both traffic and citizens' behavior. Initially, 20 professional mimes shadowed pedestrians who didn't follow crossing rules: A pedestrian running across the road would be tracked by a mime who mocked his every move. Mimes also poked fun at reckless drivers. The program was so popular that another 400 people were trained as mimes.[44] What would happen if mimes followed people in your organization around their offices for a day and mocked people who talked about sensitive information in public areas or left their computers logged on while they were away from their desks?

Social psychologist Robert Cialdini states, "Mysteries are powerful because they create a need for closure."[45] Mysteries exist wherever there are questions without obvious answers. "Why do criminals attack personal home computers?" "How was the source of the cyberattacks on Estonia discovered?" "What does the encrypted message in this week's security awareness contest say?"

Another way to maintain interest is to pose a question or puzzle that confronts people with a gap in their knowledge. As part of a security-awareness program, newsletters, an intranet site, posters, and online courses can ask challenging quiz questions. This creates two knowledge gaps: "What's the correct answer?" and "Was I right?" Providing hints and clues for more difficult security-awareness questions or contests also helps to maintain interest. Remember, the objective is to make people think about the problem, not to guess a single right answer.

Note: Don't expect people to remember much about security from an initiation day presentation when all stimuli are new.

49.5.4 Conversational Style. Use a conversational and personal style. A conversational style is useful because "people tend to pay more attention when they perceive that they're in conversation, since they're expected to follow along and hold up their end." The brain does this even when the learner is reading (e.g., the conversation is between the learner and a book, magazine, Website, or an eLearning module).[46]

First- and second-person constructions (involving "I," "we," and "you") create a feeling of conversation between the content and the reader. In five out of five studies performed in 2000, students who received material with personalized, conversational text performed better on subsequent transfer tests than those who learned with formal text.[47] Studies found that "students performed up to forty percent better on post-learning tests if the content spoke directly to the reader using a first person, conversational style rather than taking a formal tone."[48]

In addition, Dr. Roger Schank, author and expert in workplace learning, states that conversation is a form of learning by doing.[49]

49.5.5 Analogies and Examples. To better engage people with a new topic, start by highlighting things that the audience already knows. As an example, your target audience may know that the Storm worm was widely spread malicious code, but they may not know that the controlling computer changes the malicious code it sends every 30 minutes, or that the Storm worm contained new defensive techniques that shut down the efforts of researchers who were attempting to learn more about it. Use analogies and examples to tie ideas your audience knows to what you want them to learn:

- Backups are like flossing—everyone knows it's important, but few devote enough thought or energy to it.
- A dynamic IP address is like moving your house several times a day so that burglars can't find it.
- Sensitive information is like prescription medicine: it should be used only by those who need it and are authorized to have it; it should not be transferred, sold, or given to people who are not authorized to have it (this is illegal and penalties apply); it can cause damage if given to people who are not authorized to have it.[50]
- Passwords are like bubble gum; strongest when fresh; should be used by an individual and not a group; and if laying around, will create a sticky mess.[51]

49.5.6 Stories and Empathy. The number-one story commandment, according to Pixar is: "Make me care."[52] Include stories in awareness materials to increase people's attention and retention. Stories energize powers of recall and communicate priorities effectively. Stories are powerful because they provide the content missing from abstract prose. Stories captivate people. They survive the test of time, and they become part of the popular culture. Audiences are typically more receptive to someone who tells stories than to someone who lectures.

Relating a story is different from making a reasoned argument because the way the message is delivered determines how the audience will react. When the audience hears an argument, they evaluate it and usually argue back, even if only in their minds. Stories engage the audience and involve people with the idea. When an audience hears a story, they are likely to think of similar situations that they have experienced. Stories can suggest a course of action to someone who is at a decision-making point.

Stories should be about situations that are realistic and related to the lives of the learners; otherwise, stories may backfire and cause a loss of credibility. The stories should relate to situations and decisions the audience may face. Stories about hackers accessing medical records would be useful to organizations that process medical data, whereas stories about fraud or identity theft would be of interest to personnel involved in the financial industry or the accounting function of an organization.

Effective security-awareness stories are short (a few paragraphs), have two or three characters at the most, and have a singular message. Stories that show more than one point of view increase retention because the brain is tuned to learn more deeply when it is forced to make evaluations and judgments. Stories should contain a surprising element and they must be true. Lack of credibility in a story is a single point of failure.

Stories about real people and real consequences (people being praised, disciplined, or fired) are useful in presentations and courses. Sources of stories include individuals who have been with the organization for a long time and have a "corporate memory," news events, Internet special interest bulletin boards, and security personnel who attend special interest group meetings and conferences.

Organizations should collect stories about security incidents, security heroes, mistakes made, and lessons learned. Having a story collection prepared allows for quick response to trigger events, such as when an incident similar to ones from the story collection occurs at the organization. "The secret is to gather, gather, gather—and do it in advance of any pressing need. … Gather things that get a response in you. … Anything that displays or evokes energy. Storage is cheap."[53]

49.5.7 Currency. Awareness material must be fresh and current. Chef Oscar Gizelt of Delmonico's Restaurant in New York said, "Fish should smell like the tide. Once they smell like fish, it's too late." If awareness material is not changed frequently, it too begins to smell old and becomes boring.

Take advantage of circumstances. Prepare material to be ready to launch a campaign alerting people to respond quickly after a disaster or major news event. For example, after a disaster such as the tsunami in Japan there are always scams and malware that use sensational news headlines to entice potential victims to click on links to sites that contain drive-by malware.

Another idea is to prepare messages in advance for specific times of the year. November 30 is International Computer Security Day. In October, educate about scareware. During the holiday season in November and December, laptop and mobile-device theft peaks, so offer tips on securing mobile devices. For tax time, offer tips about keeping

track of your credit and information regarding fraud resolution specialists to contact for victims of fraud (e.g., remind staff that if the IRS wants more information from them, they will receive a letter, not an email. When an email claims to be from the IRS, it's best not to click on any links—and to report it to phishing@irs.gov).

Current events can be an excellent source of material and can add credibility to an awareness program. Review current news reports for events that can be used to emphasize security messages. For example, newscasters who make statements when they do not know that their microphones are live or a show on the impact of ID theft can provide good examples for use in an awareness program. Also, several Internet security and technology sites offer subscriptions to electronic security alerts and news clippings. One of the more useful newsletters reporting on information-security breaches is from INFOWAR.[54] Some organizations have established a news-hawk program, in which rewards are given to the first employee to bring in a new relevant story that can be used as part of the awareness program. This is also a good technique to gain buy-in from the end user community.

One of the best times to raise awareness is right after a breach at the organization or at a similar organization (one in the same industry, in the same location, or using the same technology). The news is full of stories about information security and data breaches.

49.5.8 Credibility. Credibility is crucial for an awareness program. The message must be clear, relevant, and appropriate to the real world. If the audience is required to use 15 different passwords as a part of day-to-day functions, prohibiting them from writing their passwords may not be as realistic as providing strategies for protecting the written list.

Show consequences. Some organizations send memos to all staff that describe specific examples of personnel who have violated policy. The memos cover a set time period (e.g., the previous quarter) and include the number of individuals, the nature of the violations, and the penalties, such as loss of Internet privileges or leave without pay, displayed in dollars, and based on the average salary.

While all messages should have a call to action, credible messages avoid fear, uncertainty, and doubt (FUD). FUD isn't the best choice for communication, and it will backfire if the material creates a scare and then doesn't offer a practical solution.

Comedian Chris Bliss explains why FUD doesn't work:

> A great piece of comedy is a verbal magic trick… there's this mental delight that's followed by the physical response of laughter, which, not coincidentally, releases endorphins in the brain. And just like that, you've been seduced into a different way of looking at something because the endorphins have brought down your defenses. This is the exact opposite of the way that anger and fear and panic, all of the flight-or-fight responses, operate. Flight-or-fight releases adrenalin, which throws our walls up sky-high. And the comedy comes along, dealing with a lot of the same areas where our defenses are the strongest—race, religion, politics, sexuality—only by approaching them through humor instead of adrenalin, we get endorphins and the alchemy of laughter turns our walls into windows, revealing a fresh and unexpected point of view.[55]

The use of appropriate humor such as the advertising technique of exaggerated consequences works because humor is a door into the serious. To spoof the popular advertisements that show a chain reaction of consequences, you could do a similar ad that says:

1. If you don't teach your employees about security awareness, they will post their vacation plans on social media sites.

2. If your employees post their vacation plans on social media sites, the bad guys will know when your employees aren't home.
3. If the bad guys know when your employees aren't home, the bad guys will clean them out.
4. If your employees' homes are cleaned out, they will come to work naked.
5. Don't let your employees come to work naked.

Another point that Chris Bliss makes is that comedy is a powerful way to communicate because it's inherently viral; people can't wait to pass along a great new joke. He adds, "But it's when you put all of these elements together—when you get the viral appeal of a great joke with a powerful punch line that's crafted from honesty and integrity, it can have a real world impact at changing a conversation."[56]

49.5.9 Social Proof. Robert B. Cialdini is considered an expert on influence. He studies and writes about the science of persuasion. He describes a common mistake that causes messages to self-destruct. It's the story of a former graduate student who had visited the Petrified Forest National Park in Arizona with his fiancée. At the park's entrance a sign stated, "Your heritage is being vandalized every day by theft losses of petrified wood of 14 tons a year, mostly a small piece at a time." The student was shocked when after reading the sign, his normally ultra-honest fiancée whispered, "We'd better get ours now."[57]

This incident inspired Cialdini and his colleagues to design an experiment where they posted two different signs. One used the concept of "negative social proof." It read, "Many past visitors have removed the petrified wood from the park, changing the natural state of the Petrified Forest." That sign also showed a picture of several visitors taking pieces of wood. The experiments placed a second sign to simply convey that stealing wood was not appropriate. The second sign said, "Please don't remove the petrified wood from the park, in order to preserve the natural state of the Petrified Forest." The accompanying image showed a lone visitor stealing a piece of wood, covered by the universal "No" symbol of a red circle with a slash through it.

The experimenters placed marked pieces of wood along various pathways and observed how the signs affected the rate of theft. They switched the signs at the entrance to the pathways, and they also used pathways with no signs posted as a control condition. The results:

- Where there was no sign, 3 percent of the wood pieces were stolen. Where the social proof sign (stating that many visitors had removed wood) was posted, the theft rate increased to 8 percent. Where the sign asked people not to steal the wood and depicted a single thief, the theft rate decreased to 1.7 percent.
- Put simply, social proof refers to our tendency to go along with the crowd and follow the most popular course of action. We do things that we see other people like us doing.
- Using negative social proof, for example, communicating the popularity of an undesirable behavior, focuses the audience on the prevalence, rather than the undesirability, of the behavior.
- The authors recommended that the park management reframe the statistics to focus attention on the number of people who respect the park's rules, which turned out to be more than 97 percent.[58]

Here's an example of how to apply this to security awareness. Todd Snapp, President of RocketReady, speaks to audiences about the human side of security. He often asks the audience to guess the most common passwords that his team of penetration testers finds in organizations where the passwords requirements include using characters from at least three sets (e.g., uppercase, lowercase, and numbers) and the passwords had to be changed every 90 days.

Audience members usually call out with guesses, but they rarely guess the answer. When Todd tells them, there is usually a collective groan and head slap as audience wonders why such a simple and retrospectively obvious answer didn't occur to them.

The answer? The season and the year: Fall2011, Winter2011, or Spring2012.

One way to present this information would be to start with "Did You Know? The most common passwords we find are…" This approach would catch attention, but it would also convey the wrong message. Despite the implied disapproval of choosing passwords that are easy to guess, the message is that such behavior is common. Putting the information out in a neutral way would act as strong social proof that many people just like the audience choose these easily guessed passwords.

A better way to present the information would be to advise people not to choose the season and year for passwords and to focus their attention on a positive behavior (e.g., use an image showing people who had chosen strong passwords speaking disapprovingly of a person in the organization who used the season and year). This makes it clear that people who use weak passwords are in the minority and have the disapproval of their co-workers. The take-away is that it's more effective to emphasize the deviance, not the popularity, of insecure behavior.[59]

49.5.10 Accessibility, Diversity, and Culture. Effective awareness materials are accessible, diverse, and culture specific. Guidelines for creating Web pages that are accessible to people with vision or hearing impairments are published by the World Wide Web Consortium (W3C). To be accessible, the Web pages should not rely on vision or sound alone to impart meaning; for example, all graphics should be labeled with text that explains the graphic, and the contrast between the text and the background should be maximized.[60] An alternative is to create and maintain two versions of an online course.

Accessible content is easy to understand. Check written program materials for ease of reading and understanding with Flesch–Kincaid readability levels or the Gunning Fog Index. Some word processors have the ability to perform two Flesch–Kincaid readability tests to indicate how difficult a passage is to read. The Flesch Reading Ease and the Flesch–Kincaid Grade Level use the same core measures (word length and sentence length), but different weighting factors. A text with a comparatively high score on the Reading Ease test should have a lower score on the Grade Level test.[61]

The Gunning Fog Index measures the readability of English writing. The index estimates the years of formal education needed to understand the text on a first reading. A Fog Index of 12 requires the reading level of a U.S. high school senior (around 18 years old). The Fog Index is commonly used to confirm that text can be read easily by the intended audience. Texts for a wide audience generally need a Fog Index less than 12. Awareness material requiring near-universal understanding should have an Index less than 8.[62]

Diversity is important to address cultural and other differences among staff. Awareness materials should suit the culture of the organization. Images of people in awareness materials should show different genders and races, with the subjects dressed similarly to the way that people within the organization dress. People relate to pictures of other

people who are similar to themselves. In a global awareness program, language and cultural differences should be addressed. Materials should avoid the use of local idioms such as (in the United States), "In a nutshell."

Awareness materials designed for use in an Islamic country might address differences in techniques used by social engineers. For example, high-pressure techniques (as are often used in America) would likely be counterproductive in an Islamic culture. Instead, social engineers are likely to concentrate on the techniques that emphasize trust and relationship building. Also, since many social engineers are most active when they expect the target company to have fewer resources, instead of noting that an attack is most likely on a Friday afternoon (as in the United States), the materials should note that attacks are most likely on Thursday afternoons and that organizations might receive a higher volume of social engineering calls during Ramadan or around the Eids when the organization may have fewer staff available.

Awareness materials designed for Japan might address threats that are targeted to the Japanese culture, where privacy, reputation, family, and a desire to stay out of trouble, such as the "It's me, It's me" fraud where a young-sounding person calls a senior and claims to be in trouble and need money to avoid a scandal, and the "One click contract" fraud, where a Website visitor to a dating or pornography site receives the message that by clicking on the previous page they have entered a contract and must now deposit money into a specified account to avoid fines or scandal. By keeping the amount affordable (e.g., the equivalent of about $100 U.S. dollars), many Japanese people simply make the deposit to avoid the potential trouble.[63]

When designing material for a company in Qatar, which has a rich history and industry of pearl diving, a memorable approach is to use a pearl to illustrate the value of data. Exhibit 49.4 shows a poster image of an oyster with a data disc in place of the pearl with the caption, "How Valuable Is Your Data?"

49.5.11 Spaced Repetition. The technique of spaced repetition was identified by Hermann Ebbinghaus about a hundred years ago. He observed that learning and memory are the strongest if you spread the repetition of information over a long period of time—for example, days, weeks, and months. He also proposed a "forgetting curve," which are like radioactive half-lives: Each review of the information to be learned increases memory in strength by about 50 percent, but immediate review does not increase memory very much because the memory hasn't decayed much.[64]

A meta-analysis in 1999 suggested that those who learn information by spaced repetition will outperform 67 percent of those who learn by mass presentation given the same number of practice episodes. This varies according to the "nature of the task being practiced, the inter-trial time interval, and the interaction between these two variables."[65]

For awareness, this means exposure to awareness once a year likely not enough to effect behavior changes. For awareness, a best practice is to follow John Medina's Brain Rules: "Repeat to remember" and "Remember to repeat."

49.5.12 Pretest and Refine Messages and Methods Before Distributing Them. Pretest and refine materials before distributing them. Pretesting provides evidence that materials are reaching your target audience with the intended message. It can also avoid embarrassing situations, such as occur after distributing a poster on which punctuation or the lack of punctuation changes the message (e.g., "Slow Work Zone" instead of "Slow, Work Zone").

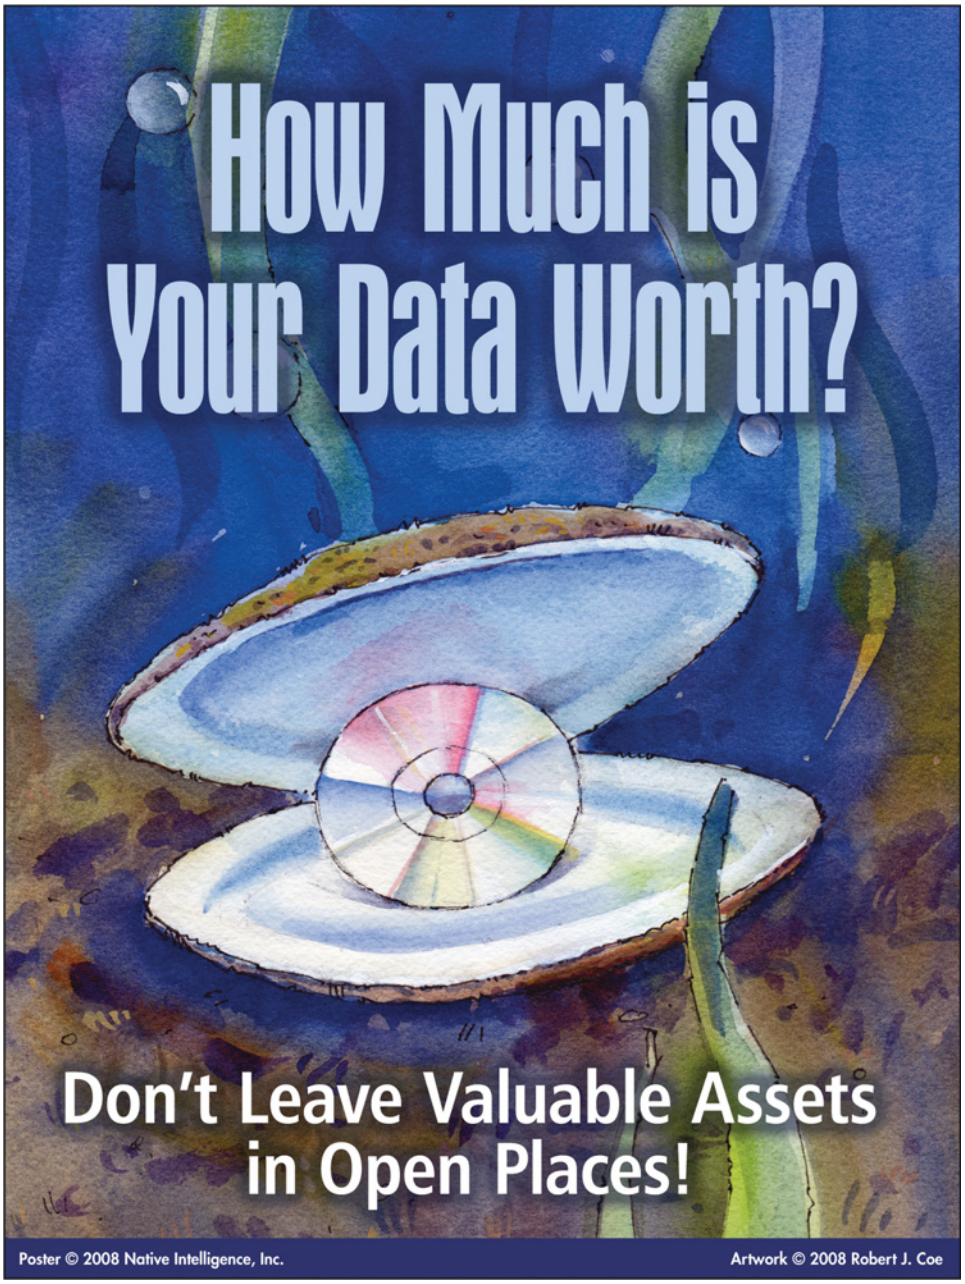


EXHIBIT 49.4 "How Valuable Is Your Data?"

Pretesting may be accomplished using focus groups, providing materials to a single unit within the organization, or through group or individual interviews. All evaluations should include a set of multiple-choice or ranking questions with one or two open-ended questions. This analysis approach facilitates data comparison and aggregation. The question set should be structured to determine the message received and the level of experience (novice, beginner, user, power user) required to understand the material. The questions should also be structured to avoid leading the respondent. A

useful approach for coordinating input from larger focus groups is Computer-Aided Consensus™, which uses a shared spreadsheet to identify hidden assumptions and divergent values among the respondents.[66]

49.6 TOOLS. When choosing tools to convey an awareness message, address these questions:

1. What tools are most appropriate for the message?
2. What methods are most likely to be credible to and accessible by the target audience?
3. Which methods (and how many methods) are feasible, considering the available budget and the time frame?
4. How often should each method and messages be delivered?

Use as many methods and tools as possible, with a consistent message, to reinforce the material and increase the likelihood that the audience will be exposed to it often enough and long enough to absorb it. Some methods are suited to daily updates, such as tips or questions of the day. Some material, such as simplified policies, frequently asked questions (FAQs), and incident reporting information will be received best if it's available on demand and just in time when needed. Newsletters are best received monthly or quarterly. Online courses can be provided monthly (short content pieces or modules), quarterly, biannually, or annually. Regular updates to content on an intranet Website, including contests, will drive repeat visits and increase awareness; stale information will lose viewers.

49.6.1 Intranet Website. An intranet Website focused on security can contain checklists (e.g., what to post on social media, how to protect mobile devices, how to manage privacy on Facebook), one-line policies (the most important concept of each policy distilled to a single line) linked to full policies, identify–react charts, and interactive technologies such as password visualizers that illustrate the relationships of your passwords and password strength meters.

Security-awareness activities that use the Internet or an intranet offer the advantages of ease of use, scalability (can be used for various audience sizes and in distributed locations), accountability (can capture use statistics, quiz or test scoring, and other metrics), accommodation of individual learning rates, and even interaction among members of a community or among students and instructors.

Websites (public or private) can be used in these ways:

- As a research tool for gathering information
- To present policies and other documents
- To post alerts
- To collect data for security-awareness surveys or incident reporting
- For self-assessments to identify at-risk security practices
- For anonymous reporting of security concerns
- For Webcasts of security conferences or presentations

49.6.2 Social Media and Crowd Sourcing. Use social marketing techniques such as Twitter and social networking, email, daily tips, questions, contests, surveys, and suggestion programs that help achieve buy-in. Your social media strategy can be used to push and pull. Use social media as a channel for delivering messages to an audience as a push. Use social media as a way to listen and learn and to create relationships as a pull. Test and learn. Focus on understanding, run experiments (for new benefits and services) and analyze the results (for audience engagement). Where traditional companies push out messages and products, these companies pull customers in. Instead of treating customers as passive targets, they treat them as active participants. Like the sun in a solar system, they create a gravitational field that pulls customers into their orbit. They go beyond customer loyalty to building customer gravity.[67]

If your organization is a nonprofit, you can use virtual volunteers to help with your awareness program. One example is Sparked.com, a micro-volunteering network, in which nonprofits post challenges to the network and volunteers respond with ideas. Challenges range from requesting user input on a logo to creating promotional materials. This brings together the talents and ideas of many to find a single solution, and gives nonprofits a way to get valuable work for free. Volunteers might design newsletters, illustrate online courses, or create Website pages. Volunteers search for opportunities by skill, interest area, development topic, or geographic region.

49.6.3 Videos and Podcasts. Videos can be delivered on DVDs, VHS tapes, CD-ROMs, or over the Web in various formats including podcasts (videos formatted to play on iPods or other portable media players). Most security-awareness videos are less than 20 minutes long. They can be used at orientation briefings and brown-bag lunches for staff where popcorn can be provided in bags preprinted with security messages. Videos are useful starting points for discussions and for briefings. They provide a consistent message throughout the organization and can be shown to staff at distributed locations, saving instructor travel time and costs. They can also be used to demonstrate cost effectively the impact of security failures, such as a fire at a data center or how sensitive data were found in the trash. Security-awareness videos are available commercially for various fees and from the U.S. Government often at no charge or for a nominal fee.

Produce awareness videos in digital format in segments that allow for updates as the environment or organizational needs change.

49.6.4 Compliance Statements. Compliance statements and policy reading sign-offs are among the most effective security-awareness offerings. After computer-based instruction, according to the Security-Awareness Index report, the most effective methods were:

- Tracking whether workers read policies or not
- Requiring a compliance statement
- Requiring full-time employees (FTEs) to read policies
- Making policies available in electronic format
- Require workers to read policies annually
- Requiring a compliance statement prior to issuing a user ID[68]

49.6.5 Sign-on Messages, Networked Screen Savers. With some systems, it is possible to add a text message to the log-on or sign-on screen. These messages should be short, to the point, and changed frequently.

Screen savers are a graphic form of communication and should be eye-catching for maximum impact. Involving a professional artist will improve message delivery. Screen savers should contain contact information for the organization's security and incident-handling functions. Animations or trivia questions and answers may make the screen saver more interesting. Screen savers should be updated periodically to keep the message fresh. Commercially produced security screen savers are available, as are screen savers that can be easily tailored to deliver a security message.

Software programs such as BroadcastIT offer a centrally managed screen saver that can show images and videos. This software updates in the background and doesn't require much space.

49.6.6 Publications. Publications, such as newsletters, brochures, pamphlets, comic books, tip sheets, identify and react sheets (documents that list signs of an incident and steps for the end user to take), whitepapers, and checklists of behaviors organized by topic, can be targeted to specific audiences. They may be security focused or may be generalized publications that contain articles on security-related events or items of interest. Newsletters should be short, one to two pages, tailored to the organization's industry or business. Newsletters should use attention-getting graphics, headlines, and white space to appeal to readers. Audience interaction can be generated by encouraging questions, answered in future newsletters, or by including contests, inviting readers to submit news items, tips, trivia, or reviews of security books or products.

49.6.7 Posters and Digital Signage. A poster series with themes or related designs can be used to highlight specific security issues. A poster should be colorful, present a single message or idea, and include a "call to action." Using a professional artist to design the posters will increase their impact. Posters should be larger than standard letter size to stand out and gain attention. They should be changed or rotated regularly and placed at eye level in multiple locations. Posters can be printed on both sides of the paper, saving paper and shipping costs for organizations with multiple locations. Signs can make a difference in behavior. In Kenya, inexpensive messages urging minibus passengers to heckle and criticize their drivers for being reckless caused a 50–60 percent reduction in insurance claims involving injury or death. The stickers had messages such as "Don't just sit there as he drives dangerously! Stand up. Speak up. Now!" and were illustrated with severed feet and legs. Drivers were given incentives to leave the stickers in place. The messages encouraging passengers to speak up were placed in a random sample of over 1,000 long-distance Kenyan minibuses. In those buses, insurance claims fell by a half to two-thirds, from 10 to less than 5 percent annually. "Results of a driver survey eight months into the intervention indicated that passenger heckling contributed to the safety"[69] improvement.

49.6.8 eLearning Courses. Web-based awareness courses are useful for geographically dispersed staff members and staff who need to take training at a time that is convenient for them (e.g., after normal work hours). Web-based courses are especially well suited for use by individuals who have diverse backgrounds and different

technology experience levels. Online courses offer the following advantages over traditional place-based, classroom training:

- Feedback—Feedback is essential to motivation and performance. Feedback is immediate, so learners do not build on early misunderstandings. Well-designed Web-based training takes cultural and personality differences into account and reassures timid trainees while allowing more confident ones to progress at a faster pace. "Why," "How," "Show me an example," and "Give me an alternative" buttons or links can be used to let learners with different needs and personalities use the course to learn in ways that are comfortable for them.
- User convenience—Web-based awareness courses are convenient for the learners because they can be taken at any time. Those with variable or hectic schedules can arrange to take the course after hours or whenever they have a convenient time in their schedule.
- Nonthreatening—Web-based courses allow users to make mistakes and learn from them in a safe, nonthreatening environment.
- Flexibility—Web-based courses are flexible and can be customized to accommodate learners with different levels of experience and different interests. By placing detailed information in subordinate, linked pages, users are able to choose between the "need-to-know" main pages and the "nice-to-know" hyperlinked pages.
- Web-based courses can reduce costs and training time. Placing updates to courses on the Web eliminates the work involved with distributing the current version and materials to multiple locations. This can be more efficient and consistent because the content has been reviewed, edited, and tested to make it clear and concise. Courseware can also be directly linked to specific organizational policies and procedures.
- Web-based courses are self-paced, so that more experienced users can race through without getting bored while novice users can ponder and explore.

A potential problem to watch for in Web-based courses is the tendency to get lost in the technology. Just because an awareness course could have three dozen animated, singing computers decorating the pages does not mean that it should. The technology must be used appropriately; bigger buildings do not make better scholars, and more impressive technology does not necessarily result in a better learning experience. A Web-based course that is overloaded with animations and graphics that do not relate to course content or that has a poorly designed user interface will lose user acceptance.

49.6.9 Classroom Training and Clickers. A study by the European Network and Information Security Agency (ENISA) found that:

> The most effective technique has been face-to-face time with staff through workshops and training sessions. Being able to put a face to a name or function is more personable and people are more receptive to messages being face-to-face. The training is mandatory. Senior management actively supports the awareness schemes, making sure that training events are at convenient times for the business and promoting them to staff. There is good attendance at sessions since missing the events results in escalation to the employee's manager. This senior management support across the business has proved to be critical to the success of the awareness program.[70]

A downside to in-person, instructor-led training is cost. Instructor-led training must have an audience that is collocated. Instructor-led training is not cost effective for organizations with a large, distributed workforce—unless they have an in-place audio/visual conference capability.

Use audience response systems (e.g., clickers) to enhance feedback and audience interaction. Audience response systems can be used to take attendance, instant polls, and multiple-choice tests. You can also use the devices to survey audience awareness of security topics and instantly show the correct answers and a response breakdown. Clickers allow online crowdsourcing for offline crowds. These devices have a fun factor and have been described as "efficient, eco-friendly and techno-tickling." Clickers allow audiences to participate in the same way as TV game-show contestants and help shy people "speak up." Professor James Katz, Director of the Center for Mobile Communication Studies at Rutgers said: "If people feel their opinions really count, they'll be happy and likely to give more opinions."[71] With the prevalence of cell phones, text messages can now be used for in-class responses.

49.6.10 People Penetration Tests and Spear Phishing Exercises. Demonstrations of penetration tests and spear phishing exercises are a good way to generate interest in security among the technical staff and power users. You can also incorporate vendors and other business partners into internal tests to evaluate and strengthen your organization's ability to respond to a cyberattack. Consulting firms may offer social engineering and penetration testing includes email phishing tests and phone calls from social engineers. The goals are to test employees' awareness of fraudulent email messages and incident response effectiveness. It must be made clear that penetration tests are intended as training exercises, not as employee evaluations.

You can also test awareness by checking the strength of passwords or simulating social engineering attacks to gauge responses. You can develop penetration exercises internally or acquire external support. Products such as Metasploit, Core Impact, and Canvass can be used to simulate a wide range of vulnerability tests, including email and Web phishing exercises. Before conducting any phishing or other vulnerability exercise, obtain permission in writing (often termed Rules of Engagement). Without permission, such activities may be considered hacking, even when performed by security personnel as part of an awareness exercise.

All testing works best when in support of defined policies and procedures. The same is true of penetration exercises. A well-designed test will include techniques for recognizing an attack as well as remedial actions that "victims" of an attack should take (e.g., reporting spear phishing attacks to organization security).

Testing can provide an estimate of the likelihood that a member of your workforce will:

- Click on an embedded hyperlink in a suspect email message,
- Enter Social Security numbers,
- Open an attachment without checking to see if it might include malware,
- Verify their network account,
- Register at an unknown site to download a whitepaper and receive a free product, or
- Download a suspicious file.

After the exercise, collect follow-up information by sending a message to recipients for feedback asking, "Why did you click or not click" or other questions to assess their thought processes.

Metrics to evaluate testing might include:

- Emails that were deleted and not read,
- Email forwarded to security personnel,
- Reports as spam,
- Emails that were read,
- Replies to email,
- Forwards to colleagues,
- "Victims" who clicked on the link, and
- "Victims" who provided personal information.

Real-world simulation exercises add to but do not replace traditional awareness techniques. Classroom instruction can explain and describe attack approaches, but exercises provide students with a "touch, feel, and experience" that expands their understanding of security attack approaches and aids in content retention. The goal of a security-awareness exercise is to make security a natural concern within the organization, campus, or university. Periodic security-awareness exercises will help minimize network downtime and maximize network performance as students become more judicious about handling emails.[72]

49.6.11 Contests and Incentive Prizes. Contests, incentive prizes, rewards, and giveaways help achieve buy-in. People like to win and most love a good contest. A contest can be a simple prize draw or a competition with rules for entry and criteria for winning.

Contests can be used to ask questions, collect data, conduct research, inspire ideas, or drive traffic to your security Website. The 3 Ps of contests are:

- **Planning.** Decide on the goal of your contest, then create a theme. Establish clear rules, including entry procedures and criteria for judging competition entries. You may need to consider your country or state-specific regulations (e.g., some contests and competitions may require a permit, if open to the public and the competition is a random chance draw).
- **Prizes.** Prizes can be anything from bragging rights to security-themed DVDs (such as *Catch Me If You Can* or *Swordfish*) or books (such as *The Cuckoos' Egg* by Clifford Stoll, *Kingpin* by Kevin Poulsen, or *The Lure* by Steve Schroeder) to cash awards, lunch with a senior executive, or time off. At presentations, speakers can tape prizes or awards under seats in the front row to encourage people to come early and sit up front.

  Prizes don't have to be expensive to be valuable to your audience. Time off, lunch with the boss, gift certificates, shredders, security-themed T-shirts, mugs, certificates, and trophies all work. Shiny prizes, such as the latest technology (e.g., an iPad) or money, have mass appeal. You may want to poll your audience to find out what would be valuable to them. The prize might be intangible (e.g., an honor, such as announcing the winner in the organization newsletter).

- **Promotion.** Announce your contest with email, posts on the security intranet site, and posters. Also, if it's a public contest, consider tweeting it, enlisting coworkers to help spread the word, putting links to the contest on the organization's Facebook Page, and writing press releases.

If you run the contest on social media, such as Facebook, Twitter, or Google+, use applications (many are available) to help administer the contest. To win, people may simply have to follow, retweet, or answer trivia questions. This type of campaign is often successful because of the ease of entry. Your security "brand" will benefit from the increased engagement.

Here are some contest ideas:

- **What's That Number?** Post a number that relates to security and have people guess how the number relates to security. The number might be the number of password reset requests the helpdesk receives in a week, the number of malware-infested sites blocked by the corporate firewall, or the number of records exposed or dollars lost as a result of a breach experienced by a company in your industry.
- **Catch the Red Team**. A red team is a group of individuals assigned to test the security of an organization. Staff are told that a red team will be testing security (e.g., making social engineering calls). Staff members who catch the red team and report the potential security violations win a prize. Often, these contests result in identifying security vulnerabilities and sometimes in catching intrusion attempts by cybercriminals and not just the attempts of the red team.
- **Nooo Face!** A security-awareness video or photo contest, such as the Annual Security Video Contest held by Educause or Trend Micro's "Nooo! Face" Contest, provide awards for photos that capture the feeling one gets when they realize they are a cyber, victim: precious data has vanished, destroyed, or been taken by an attack.
- **Awareness Materials Contests**. Award prizes and recognition for awareness materials, such as the annual contests held by the Federal Information Systems Security Educators' Association (FISSEA) and the International Information Systems Security Certification Consortium, Inc., ((ISC)^2)'s CyberExchange. FISSEA's contest has categories: Awareness Posters, Motivational Items (trinkets—pens, stress relief items, T-shirts, etc.), Awareness Websites, Awareness Newsletters, and Role-Based Training & Education. The CyberExchange accepts posters, presentations, best practices, flyers, white papers, and more.
- **Security Song, Jingle, and Verse Contests**. Contests could be for the best security haiku or six-word security stories. The six-word stories are based on the challenge issued to Ernest Hemingway to write an entire story in six words. His story: "For sale: baby shoes, never worn." A security-related six-word story might be "I never checked my offsite backup…"
- **Top Ten Lists**. Award a prize for the best (funniest) Security Top Ten list. Examples are the "Top Ten Places Not to Hide Your Password" (such as written with a permanent marker on a light bulb in the office lamp, on a white board, as a tattoo) and "Top Ten Security Headlines We'll Never See" such as, "White House Painted Purple to Confuse Terrorists" or "Courts Close Due to Lack of Lawsuits over Security Breaches."

- **Security Stories**. Invite people to share their security stories—for example, how a person was affected by identity theft, or how someone refused to share personal data when it wasn't necessary to do so, such as when a healthcare provider's form asks for a Social Security number. The emotional content makes stories prime material for sharing.
- **Security Trivia**. Ask security-related questions, such as:
  - "What color is Whitfield Diffie's hair?" (or, "Who is Whitfield Diffie and why is he important to security?")
  - "What is the name of the 1983 movie where Mathew Broderick played a young hacker who gained access to a government nuclear war simulator?" (*War Games*)
    Or, for more technically advanced audiences,
  - "What type of attack against database-driven applications involves the intruder manipulating a site's Web-based interfaces to force the database to execute undesirable code?" (SQL injection)
  - "What hardware protocol caused the vulnerability where a Firewire device, when plugged in, can overwrite anywhere in memory?" (DMA or Direct Memory Access)
- **Security Fact or Fiction Contests**. Contestants must decide if statements are true or not. For example: "In fiscal year 2011, the Electronic Crimes Special Agent program processed 1,066 terabytes of data on 8,525 units" (fact).[73]

Contests can boost morale, motivate people, and contribute to team spirit. Vince Lombardi, former head coach of the Green Bay Packers, understood this. He once said, "Winning isn't everything. It's the only thing." After criticizing him for this statement, some of his critics put together a new kind of baseball league for children in a Texas community: "It was like the Little League—the same ball, same bat, same number of innings, same playing field—everything was the same except that they didn't keep score. The idea was that there wouldn't be any losers because nobody would know who won." The game lasted one and a half innings. After that "the kids went across the street to play sand-lot ball where they could keep score."[74]

49.6.12 Awards and Recognition. Rewarding good security behaviors contributes to good security. Security is part of everyone's job and often management believes no special recognition or incentives should be provided. This approach does not work well because in a poorly managed organization, security tends to be outside the normal business process (i.e., security measures are often viewed as an impediment to getting a job done).

Security is a special concern that must be emphasized if assets are to be adequately protected. Security should be integrated with performance appraisals. "Personnel become motivated to actively support information security and privacy initiatives when they know that their job advancement, compensation, and benefits will be impacted. If this does not exist, then an organization is destined to depend only upon technology for information security assurance."[75]

Awards can be given for extraordinary security behaviors, participation in security events, achieving a security certification, or providing a security service such as speaking to local groups about cybersecurity. Awards work best if they are publicized, support desired behaviors, and are immediate (close to the act for which the award is being given).

49.6.13 Human Libraries. A Human Library is a technique to promote dialogue, information interchange, reduce prejudices, and encourage understanding. A Human Library consists of a group of individuals ("books") who have agreed to share their knowledge (i.e., the information that's in their head) with others. Living Books are people you have recruited because they have experiences of interest to the audience. A Human Library can be established as a single event with defined start and end times or an ongoing activity (long-term resource), where the Living Books come and go much as books are checked in and out from a conventional library; it may be established in a single physical location where the "books" are available for a fixed time frame, or it may be a virtual library where the books may be checked out by accessing a database.[76] For security awareness, a Human Library can serve as a resource where members of the workforce can learn about security from people who understand their work environment their specific policies and procedures.

Living Books should be volunteers who are recruited with care to ensure that they are committed and willing to talk with strangers about important and sometimes very personal issues. Interview book candidates to ensure the quality of books. Ask the book about its title (subject area) and motivation to be a book. This is to ensure that books are focused on supporting awareness. A reader can safely ask any question without fear of ridicule. A Human Library provides an opportunity to ask the information security questions you always wanted to ask, but were afraid that asking would make you appear naïve.

The best sellers are defined as the books that have the most requests for loans. For metrics, ask books, readers, and librarians for their comments on their Human Library experience. Ask the books if they would be a book again. Ask if people felt that they benefited from the library. Ask the books if they learned anything from the readers.

The experiences that might increase security awareness in living books include:

- Victim of identity theft
- Computer gaming addict
- Computer Incident Response Team member
- Penetration tester
- Social engineer
- Digital forensics expert
- Hacker
- Ethical hacker
- Helpdesk staff member
- Biometric expert
- Reformed cyberbully
- Someone who lost their job as a result of something posted on the Internet
- Information system security officer
- Senior executive responsible for security policy
- Privacy expert
- Electronic Frontier Foundation member
- Information-security blogger

- HIPAA expert
- Malware researcher
- Computer programmer

49.6.14 Volunteer Activities. Strengthen your security-minded workforce through volunteer programs. Volunteerism increases engagement. Data from the Center for Talent Innovation (CTI) shows that the vast majority of college graduates want to amplify their commitment to good causes through their employer.[77] According to the Deloitte Volunteer Impact Survey,[78] Generation Ys who frequently participate in their company's volunteer activities are more likely to be very proud to work for their company, feel very loyal, and be very satisfied with the progression of their careers. These sentiments hold true across generational cohorts—91 percent of Gen X women and 76 percent of Gen X men, and 90 percent of female and 79 percent of male Baby Boomers, feel it is important to contribute to their community or the wider world through work.[79]

There are many programs that allow individuals with security interests to give back to the community:

- Attending the FBI Citizens Academy, which shows how dedicated our FBI is to protecting our freedom. Most individuals not exposed to computer crime get a rude awakening of just how bad cybercrime is, especially when they see in real-time the innocent images of victims of pedophile activity in their own neighborhood.
- Individuals with CISSP certification can join the (ISC)^2 Safe and Secure Online Program to teach children how to be cyberaware.
- Volunteering to be a living book at a Human Library event.
- Participating in National Teach-In day to promote cyberawareness for elementary students.
- Cyberawareness training for seniors can be arranged through libraries, religious organizations, and other groups.

49.6.15 Inspections and Audits. Inspections and audits raise awareness among the staff being reviewed, at least for the duration of the inspection. Audits and inspections are typically viewed as negative events. However, there are approaches that can turn an audit/inspection from a negative to a positive experience. Using a technique called "security by wandering around" (SBWA), a security staff member tours the work area, identifies staff members doing something correctly, and leaves certificates of congratulations, thank-you notes, or trinkets on their desks. One audit technique is to treat each encounter with a staff member as a training opportunity, explanations are provided as to "why" a policy is important instead of just rating compliance as pass or fail. Security personnel might periodically demonstrate social engineering by attempting to smooth-talk users into providing their passwords. The number of people who fall for the scheme might be used as an example for the next awareness session.

49.7 EVALUATION AND METRICS. Security consultant Gary Hinson compared security to the brakes on a car. The brakes slow you down, but they also make it possible for you to go a lot faster. A good metrics program takes time to set up, but once you have it set up and working well, it can save you time in the long run by making

your program more effective. Metrics aid in decision making. Without a solid metrics program, it is hard to know whether the program is effective, or if the organization should spend more money on doing the same thing, or if the resources would be better used elsewhere.

49.7.1 Baseline. As with any tool, it is important to know how to use metrics. Metrics are best used when they compare measurements over time to a baseline. Defining appropriate metrics and immediately gathering data on the current state of those metrics is an essential first phase for all future evaluation of security-awareness programs.

49.7.2 Metrics. Security metrics are evolving, and different organizations have put forth different tools, different guidance, and different frameworks for evaluating IT security. For example, organizations may use metrics based on such standards as:

- COBIT[80]
- FISMA[81]
- FITSAF[82]
- GLBA[83]
- HIPAA[84]
- ISO 27002:2005[85]
- NIST SP 800-55[86]
- PCI DSS[87]
- SANS[88]
- SOX[89]

or combinations of such guidance.

Few organizations currently use security metrics or even a common vocabulary. Also, many of the tools and guides for measuring IT security metrics only consider security awareness as a small portion of overall security program metrics. Metrics regarding security-awareness programs are high level and not specific. It is easy to collect quantitative measures of data, such as the number of virus infections, server patches performed, or program costs. It is difficult to measure behavioral change.

A commonly used metric is the number of people who participated in awareness orientations and refreshers. This figure can be determined through attendance sheets, course registrations, or completion notifications for online courses, and signed user "acceptance of responsibilities" statements. Another common metric is seat time—how long this person spent in front of a computer, clicking through the screens, soaking up the knowledge that was there. Attendance and seat time may indicate that a program is *not* effective, but they are not the best measures of awareness program *effectiveness*. Better measurements focus on the end users and measure behaviors that are a part of normal business operations, including user perceptions, activities, and response to anomalous occurrences.

An effective, measured awareness program can ensure that the workforce serves as a staff firewall, protecting the organization's information assets, and ensuring that there exists a gold-standard or best-in-class security environment. It also provides

decision makers who are allocating resources with assurances that the security-awareness programs are cost-effective control measures. Also, using objectives such as "ensuring that 100 percent of employees take the awareness course" is a statement regarding the program process, not impact, and should be avoided. When it comes to making a decision between a business case supported by hard numbers and one based on subjective feelings and unsupported statements such as "workforce training is the most cost-effective security control," the hard numbers will win.

Once the performance measures are adopted, the next step in building the business case solution is to establish at least one metric for each of the supporting goals and objectives. These metrics may be hard numbers (e.g., number of security incidents), estimates (e.g., the average cost per security incident), or the results of testing (e.g., annual testing of a sample of the workforce). Examples of metrics that might be used for the four awareness goals presented above include:

- Goal: The awareness program will improve employees' ability to recognize and report potential threats and vulnerabilities.
  - Metric: The number of security events as measured by the number of incident reports.
  - What to expect: When an awareness program is first introduced, the expectation is that the number of reported incidents will increase as employees become more aware of potential threats, vulnerabilities, and the need to report. Over time, the number of reported events should stabilize and decline as the security environment is strengthened.
- Goal: The awareness program will improve the level of compliance with company physical and computer security controls.
  - Metric: The number of sanctions of individuals for failure to comply with security policy.
  - Metric: The number of incidents resulting from employee action or inaction as determined through incident analysis.
  - What to expect: The number of employees sanctioned for compliance failures should decrease as the awareness program reaches more individuals. Similarly, the awareness program should make individuals more diligent in performing security-related responsibilities, thus resulting in a reduced number of incidents.
- Goal: The awareness program will reduce the occurrence of security failures resulting from employee action or inaction.
  - Metric: The number of incidents resulting from employee action or inaction as determined through incident analysis.
  - Metric: The number of security incidents resulting from employee actions or inactions declines over time as individuals recognize the importance of fulfilling their security responsibilities.
- Goal: The awareness program will reduce the severity of the security incidents that do occur.
  - Metric: The cost per incident as determined through incident analysis.
  - What to expect: The cost per incident decreases as individuals react more quickly to identify a potential security incident and take action to mitigate its impact.

Two other issues should be addressed relative to an awareness program's goals and metrics. First, an awareness program seeks to make a *behavioral* change in the workforce. Behavioral change takes time, so the metrics should have an established time frame where management should expect to see performance improvement. The time frame should be realistic (e.g., quarterly, semi-annual, annual), avoid immediate impacts (e.g., "After a course on viruses, the number of virus incidents should decrease within a week") and avoid time frames that are too long (e.g., "The number of virus incidents should decrease over the next five years").

Second, solicit and use employee feedback. Employee feedback should be analyzed because it provides indicators of the level of security awareness and the importance employees ascribe to information security. Analysis of employee feedback could be used to generate other performance metrics for an awareness program.

Automated metrics are useful to help measure changes in behavior.

- Filtering software that monitors content (e.g., 9- or 16-digit strings of numbers for Social Security numbers and credit card account numbers) or specific words can indicate how often people try to email this information;
- Web statistics and firewall monitoring software can indicate how often people visit or attempt to visit specific Websites, such as the security intranet page;
- Helpdesk call log summaries can identify problems related to security issues; and
- Performance appraisals, participation in contests, and results of quizzes are indicators of interest in security-related behaviors.

49.8 CONCLUDING REMARKS. Awareness among an organization's staff is vital to maintaining the integrity of data and systems. Although organizations often view computer security as a technological problem and use sophisticated hardware and software solutions to control access, detect potential security incidents, and prevent fraud, the reality is that computer security is as much a people problem as a technological problem. End users are closer to potential problems; therefore, they need to be aware of potential risks, threats, vulnerabilities, and their own security responsibilities.

People are major contributors to the IT security problem, and they are also crucial to its solution. People are perceptive and adaptive, and if trained and motivated to be aware, they can be the strongest and most effective security countermeasure. Individuals often are the first to detect security incidents. The actions they take or fail to take determine the level of damage. An aware workforce often can compensate for deficiencies in technical controls. The intent of the awareness program is to make recognition of and reaction to security threats a reflexive behavior.

Awareness takes time. It also requires the organization to have an in-place information security policy, the support of the senior-level managers, and clear goals and plans for achieving awareness. The importance of establishing measurable goals cannot be overestimated and is critical to obtaining support and funding.

The goal of an awareness program is often to change attitudes and behaviors that may be embedded in long-term procedures or habits. To effect awareness, the program must appeal to the audience and be tailored to the workforce and to the technology of the organization. The primary message of a security-awareness program should be that security is everyone's responsibility. Actions taken by end users make a significant difference; thus, a well-trained and motivated workforce is a critical and necessary security control.

49.9 GLOSSARY

Awareness—being conscious of what is going on around one; security awareness specifically focuses attention on security. Security awareness is the individual's understanding that security is important and that everyone has a role in ensuring the security of information and information technology.

Awareness campaign—the activities associated with conveying a specific awareness message (e.g., telling people "log off when away from your computer").

Awareness program—the planned implementation and control of a mix of awareness activities over a period of time, with measurable goals and multiple topics. An awareness program may encompass several campaigns.

Basics and literacy—a transitional stage between awareness and role-based training.

Education—the process of integrating all security skills and competencies into a common body of knowledge, adding a multidisciplinary study of concepts, issues, and principles.

End user (also computer user or user)—any person who uses an information system.

Focus group—a small group of end users (or of individuals from the target audience) who review and discuss awareness activities, courses, products, and the like, often under the guidance of an awareness material developer or training specialist.

FUD factor—the effects of fear, uncertainty, and doubt (FUD).

Gamification—the concept of applying game-design thinking to nongame applications to make them more fun and engaging.

Malicious code (also malware)—hardware, software, or firmware that is intentionally included in a system for an unauthorized purpose (e.g., a Trojan horse).

Orientation briefing—a presentation that provides new employees, contractors, and the like with basic security information and information on the organization's security policies and programs. Usually these presentations are conducted on arrival or shortly thereafter.

Refresher—an awareness activity, such as a briefing, intended to reinforce and update awareness of security controls and policies and to remind individuals of their security responsibilities.

Role-based training—the process of producing relevant and needed security skills and competency. Security awareness is the "what." Role-based training is the "how."

Safe failure—the opportunity to learn from mistakes privately, such as with a computer simulation or course.

Social engineering—social methods (e.g., threats, misrepresentations) that deceive a victim so that the victim does what the attacker wants him or her to do. Often the goal is to get the victim to provide private or sensitive information, such as account numbers or passwords. An example of a social engineering attack is the use of "phishing" emails.

Social marketing—an approach to security awareness using attraction and persuasion techniques designed to encourage a group of people to alter old ideas,

understand and accept new ideas, and value their new awareness enough to change attitudes and take positive actions to improve IT security.

Target audience—a specified audience or demographic group for which a security-awareness message is designed.

Threat—anything that can potentially harm a system or its associated assets (hardware, software, data, operations). Threats may be man-made or natural occurrences. Awareness programs do little to address threats; instead, they seek to reduce vulnerabilities.

Vulnerability—a weakness in automated system security procedures, administrative controls, physical layout, internal controls, and so forth, which could be exploited by a threat to gain unauthorized access to information or disrupt critical processing. A goal of security-awareness programs is to reduce behavior-related vulnerabilities.

49.10 NOTES

1. Katy Stech, "Burglary Triggers Medical Records Firm's Collapse," *The Wall Street Journal*, March 12, 2012, http://blogs.wsj.com/bankruptcy/2012/03/12/burglary-triggers-medical-records-firm%E2%80%99s-collapse
2. Ellis Smith, "BlueCross Burglary Could be Chattanooga's Costliest Caper," *Times Free Press*, March 13, 2012, http://timesfreepress.com/news/2012/mar/13/bluecross-burglary-could-be-citys-costliest
3. Alex Pham, "Sony Expects Much Wider Annual Loss," *Los Angeles Times*, February 3, 2012, http://articles.latimes.com/2012/feb/03/business/la-fi-ct-sony-earns-20120203
4. Tony Busseri, "It's Time to Take Cybersecurity Seriously," *Wired*, March 12, 2012, www.wired.com/threatlevel/2012/03/opinion-busseri-cybersecurity
5. "Verizon 2012 Data Breach Investigations Report," p. 2, www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012-ebk_en_xg.pdf (url inactive).
6. Mark Wilson, Kevin Stine, and Pauline Bowen, "Information Security Training Requirements: A Role- and Performance-Based Model (Draft)," NIST Special Publication 800-16 Rev. 1 (Draft), March 2009, http://csrc.nist.gov/publications/drafts/800-16-rev1/Draft-SP800-16-Rev1.pdf
7. Wilson et al., "Information Security Training Requirements," NIST Special Publication 800-16 Rev. 1 (Draft)
8. Wilson et al., "Information Security Training Requirements," NIST Special Publication 800-16 Rev. 1 (Draft)
9. Ira Winkler, *Zen and the Art of Information Security* (Rockland: Syngress, 2007).
10. Mary Kirwan, "Education May Not Be Enough to Ensure Compliance," *The Globe and Mail*, March 31, 2012, http://v1.theglobeandmail.com/servlet/story/RTGAM.20070911.WBsecurityblog20070911120000/WBStory/WBsecurityblog
11. WebCPA staff, "IRS Security Still Lax," *Accounting Today*, Washington, D.C., August 6, 2007, www.accountingtoday.com/news/25009-1.html
12. Winkler, *Zen and the Art of Information Security*.
13. "Many Young Workers Are IT Rule-Breakers," *Harvard Business Review* | The Daily Stat, January 10, 2012, http://web.hbr.org/email/archive/dailystat.php?date=011012

14. Jeffrey M. Stanton, Kathryn R. Stam, Paul Mastrangelo, and Jeffrey Jolton, "Analysis of End User Security Behaviors," *Computers & Security* 24, no. 2 (March 2005): 124–133, www.sciencedirect.com/science/article/pii/S0167404804001841
15. Thomas Steward and Anand Raman, "Lessons from Toyota's Long Drive, An Interview with Katsuaki Watanabe," *Harvard Business Review*, July 2007, http://hbr.org/2007/07/lessons-from-toyotas-long-drive/ar/1
16. John Caddell, "How to Bounce Back from a Big Mistake," *99U*, approximately November 2011, http://99u.com/articles/7089/how-to-bounce-back-from-a-big-mistake
17. D. Verton, "Federal Agency Faces Judicial Ultimatum," *Computerworld*, 2002, www.computerworld.com/securitytopics/security/story/0,10801,69937,00.html
18. María Cristina Caballero, "Academic Turns City into a Social Experiment: Mayor Mockus of Bogotá and His Spectacularly Applied Theory," *Harvard University Gazette*, March 11, 2004, www.news.harvard.edu/gazette/2004/03.11/01-mockus.html
19. M. E. Kabay, "The Net Present Value of Information Security: A Paradigm Shift for INFOSEC and E-commerce," 2006, www.mekabay.com/infosecmgmt/npvsec.pdf
20. Additional note: See also Chapter 50, "Using Social Psychology to Implement Security Policies."
21. The Federal Information Security Management Act of 2002, Section 3544, "Federal Agency Responsibilities," http://csrc.nist.gov/drivers/documents/FISMA-final.pdf
22. See Chapter 64, "U.S. Legal and Regulatory Security Issues"
23. Dan Heath and Chip Heath. *Made to Stick* (New York: Random House, 2007).
24. C. Heath and D. Heath, "The Curse of Knowledge." *Harvard Business Review*, December 2006, http://hbr.org/2006/12/the-curse-of-knowledge/ar/1
25. Donna Mattick, personal correspondence, March 2012.
26. John Tierney and Roy Baumeister, *Willpower: Rediscovering the Greatest Human Strength* (The Penguin Press HC, 2011).
27. V. Basili, G. Caldiera, and H. D. Rombach, "The Goal Question Metric Approach," *Encyclopedia of Software Engineering* (New York: John Wiley & Sons, 1994), 528–532 www.cs.umd.edu/~basili/publications/technical/T87.pdf (url inactive).
28. Mark Wilson and Joan Hash, "Building an Information Technology Security Awareness and Training Program," NIST SP 800-50, October 2003, http://csrc.nist.gov/publications/nistpubs/800-50/NIST-SP800-50.pdf
29. K. Sanders, "Bears in Your Backyard? Guidelines for Bear-Proofing Your Property, and Living *with* Bears," 2000, www.yellowstone-bearman.com/B_housesafe.html
30. Rebecca Herold, *Managing an Information Security and Privacy Awareness and Training Program*, 2nd ed. (CRC Press, 2011).
31. K. Hall and the SE SIG Steering Committee, "A System for Gaining Management Support for Your Safeguards and Security Awareness Program," 2002, www.orau.gov/se/Products/SE%20SIG%20Gaining%20Mgmt%20Support.doc
32. Wilson et al., "Information Security Training Requirements," NIST Special Publication 800-16 Rev. 1 (Draft).
33. de Saint-Exupery, Antoine. *The Wisdom of the Sands* (Amereon Ltd., October 2003).
34. John Medina, *Brain Rules* (Pear Press, 2008).

35. Medina, *Brain Rules*.
36. Reif Larson, "This Chart Is a Lonely Hunter: The Narrative Eros of the Infographic," *The Millions*, February 7, 2012, www.themillions.com/2012/02/this-chart-is-a-lonely-hunter-the-narrative-eros-of-the-infographic.html
37. Larson, "This Chart Is a Lonely Hunter."
38. Brenda Oldfield, "Game-Changing Technologies for Cybersecurity Awareness and Training" (presentation, FISSEA Conference, Gaithersburg, MD, March 29, 2012).
39. Anderson, Chris. "How Web Video Powers Global Innovation," Technology, Entertainment, Design (TED) Talks, TEDGlobal 2010, Filmed July 2010, Posted September 2010; www.ted.com/talks/chris_anderson_how_web_video_powers_global_innovation.html
40. Anderson, "How Web Video Powers Global Innovation."
41. Chip Heath and Dan Heath, *The Myth of the Garage* (Crown Business, 2011).
42. Winifred Gallagher, *New: Understanding Our Need for Novelty and Change* (Penguin Press HC: 2011).
43. Maria Popova, "Network: The Secret Life of Your Personal Data," January 10, 2012; www.brainpickings.org/index.php/2012/01/10/network-michael-rigley/ contains embedded video, "Network," by Michael Rigley
44. Caballero, "Academic Turns City into a Social Experiment."
45. Heath and Heath, *Made to Stick*.
46. Elisabeth Freeman and Eric Freeman, *Head First HTML with CSS & XHTML* (O'Reilly Media, 2005).
47. Ruth Colvin Clark and Richard E. Mayer, *eLearning and the Science of Instruction*. (San Francisco: Pfeiffer, 2003).
48. Freeman and Freeman, *Head First HTML.*
49. Roger, C. Schank, *Lessons in Learning, e-Learning, and Training: Perspectives and Guidance for the Enlightened Trainer* (Pfeiffer, 2005).
50. Native Intelligence, Poster 153A, www.nativeintelligence.com/ni-posters/posters.asp
51. Native Intelligence, Poster 115, www.nativeintelligence.com/ni-posters/posters.asp
52. Andrew Stanton, "The Clues to a Great Story," Technology, Entertainment, Design (TED)Talks, TED2012, Filmed February 2012, Posted March 2012; www.ted.com/talks/andrew_stanton_the_clues_to_a_great_story.html
53. Gerald Weinberg, *Weinberg on Writing: The Fieldstone Method* (Dorset House, 2005).
54. INFOWAR: The Nexus of Technology and Security in Cyberspace, www.infowar.com
55. Bliss, Chris. "Comedy Is Translation," Technology, Entertainment, Design (TED)Talks, TEDX Talk, www.ted.com/talks/chris_bliss_comedy_is_translation.html
56. Bliss, "Comedy Is Translation."
57. Robert B. Cialdini, Noah J. Goldstein, and Steve Martin, *Yes!: 50 Scientifically Proven Ways to Be Persuasive* (Free Press, 2009).
58. R. B. Cialdini, *Influence: The Psychology of Persuasion* (HarperBusiness, 2006).

59. M. E. Kabay, "Applying the Science of Persuasion to Security Awareness," *Network World Security Strategies Newsletters*, 2009, www.mekabay.com/nwss/725_applying_the_science_of_persuasion_to_security_awareness.pdf
60. W3C, "Accessibility," 2012, www.w3.org/standards/webdesign/accessibility
61. Rudolph Flesch, *The Art of Readable Writing* (Wiley, 1994).
62. Robert Gunning, *The Technique of Clear Writing* (McGraw-Hill, 1952).
63. Carnegie Mellon University Information Networking Institute (INI) "Confronting One-Click Fraud in Japan," www.ini.cmu.edu/news/features/one-click.html
64. Gwern.net, "Spaced Repetition," March 12, 2012, www.gwern.net/Spaced%20repetition
65. J. J. Donovan and D. J. Radosevich, "A Meta-analytic Review of the Distribution of Practice Effect: Now You See It, Now You Don't." *Journal of Applied Psychology* 84, no. 5 (1999): 795–805.
66. M. E. Kabay, "CAC: Computer-Aided Consensus™," 2009, www.mekabay.com/methodology/cac_ppt.zip
67. Mark Bonchek, "How Top Brands Pull Customers into Orbit," *Harvard Business Review* | Blogs, March 5, 2012, http://blogs.hbr.org/cs/2012/03/how_top_brands_pull_customers.html
68. PentaSafe. "Security Awareness Index Report: Worldwide State of Security Awareness," 2002.
69. James Habyarimana and William Jack, "Heckle and Chide: Results of a Randomized Road Safety Intervention in Kenya," Center for Global Development, Working Paper No. 169, April 2009, http://www.cgdev.org/files/1421541_file_Habyarimana_Jack_Heckle_FINAL.pdf
70. The European Network and Information Security Agency (ENISA), "Information Security Awareness Initiatives: Current Practice and the Measurement of Success," www.enisa.europa.eu/doc/pdf/deliverables/enisa_measuring_awareness.pdf
71. Jan Hoffman, "Speak Up? Raise Your Hand? That May No Longer Be Necessary," *The New York Times*, March 30, 2012, www.nytimes.com/2012/03/31/us/clickers-offer-instant-interactions-in-more-venues.html
72. Aaron Ferguson, "Fostering E-Mail Security Awareness: The West Point Carronade," EDUCAUSE, 2004, www.educause.edu/ir/library/pdf/EQM0517.pdf
73. United States Secret Service, "Fiscal Year 2011 Annual Report," 2011, p. 40, www.secretservice.gov/USSS_FY2011AR.pdf
74. C. Coonradt, *The Game of Work* (Park City, 1997).
75. Herold, *Managing an Information Security and Privacy Awareness and Training Program*.
76. Human Library Website, 2012, http://humanlibrary.org
77. Sylvia Ann Hewlett, "Strengthen Your Workforce Through Volunteer Programs," *Harvard Business Review* | Blogs, March 5, 2012, http://blogs.hbr.org/hbr/hewlett/2012/03/strengthen_your_workforce_thro.html
78. Deloitte Website, "2011 Deloitte Volunteer Impact Survey," 2011 www.deloitte.com/view/en_US/us/About/Community-Involvement/volunteerism/impact-day/f98eec97e6650310VgnVCM2000001b56f00aRCRD.htm
79. Hewlett, Sylvia Ann. "Strengthen Your Workforce Through Volunteer Programs," 2012; http://blogs.hbr.org/hbr/hewlett/2012/03/strengthen_your_workforce_thro.html

80. Information Systems Audit and Control Association (ISACA) "COBIT 5: A Business Framework for the Governance and Management of Enterprise IT," ISACA Website, 2012, www.isaca.org/COBIT/Pages/default.aspx
81. National Institute of Standards and Technology (NIST), "FISMA FAQs," NIST Website, 2012, http://csrc.nist.gov/groups/SMA/fisma/faqs.html
82. National Institute of Standards and Technology (NIST), "The Federal Information Technology Security Assessment Framework," 2000, http://csrc.nist.gov/drivers/documents/Federal-IT-Security-Assessment-Framework.pdf
83. Bureau of Consumer Protection, "Gramm-Leach-Bliley Act," 2012, http://business.ftc.gov/privacy-and-security/gramm-leach-bliley-act
84. U.S. Department of Health & Human Services, "Health Information Privacy," 2012, www.hhs.gov/ocr/privacy
85. ISO/IEC 27002:2005, "Information Technology—Security Techniques—Code of Practice for Information Security Management," 2005, www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=50297
86. E. Chew, M. Swanson, K. Stine, N. Bartol, A. Brown, and W. Robinson, "Performance Measurement Guide for Information Security," NIST SP-800-55, rev. 1, 2008, http://csrc.nist.gov/publications/nistpubs/800-55-Rev1/SP800-55-rev1.pdf
87. VISA Data Security Bulletin, "Visa PCI DSS Compliance Validation Framework," 2008, < http://usa.visa.com/download/merchants/cisp-bulletin-visa-pci-dss-framework-111808.pdf
88. SANS (2012). "Standards," SANS | Reading Room, 2012 www.sans.org/reading_room/whitepapers/standards
89. U.S. Securities and Exchange Commission, "The Laws That Govern the Securities Industry," U.S. SEC Website, 2012, www.sec.gov/about/laws.shtml

CHAPTER 50

USING SOCIAL PSYCHOLOGY TO IMPLEMENT SECURITY POLICIES

M. E. Kabay, Bridgitt Robertson, Mani Akella, and D. T. Lang

50.1 INTRODUCTION 50·1

50.2 RATIONALITY IS NOT ENOUGH 50·2

- 50.2.1 Schema 50·3
- 50.2.2 Theories of Personality 50·4
- 50.2.3 Explanations of Behavior 50·7
- 50.2.4 Errors of Attribution 50·7
- 50.2.5 Intercultural Differences 50·10
- 50.2.6 Framing Reality 50·11
- 50.2.7 Getting Your Security Policies Across 50·12
- 50.2.8 Reward versus Punishment 50·13

50.3 BELIEFS AND ATTITUDES 50·13

- 50.3.1 Beliefs 50·14
- 50.3.2 Attitudes 50·14
- 50.3.3 Changing Attitudes toward Security 50·14

50.4 ENCOURAGING INITIATIVE 50·16

- 50.4.1 Prosocial Behavior 50·16
- 50.4.2 Conformity, Compliance, and Obedience 50·17

50.5 GROUP BEHAVIOR 50·20

- 50.5.1 Social Arousal 50·20
- 50.5.2 Locus of Control 50·20
- 50.5.3 Group Polarization 50·20
- 50.5.4 Groupthink 50·20

50.6 TECHNOLOGICAL GENERATION GAPS 50·21

50.7 SUMMARY OF RECOMMENDATIONS 50·22

50.8 FURTHER READING 50·24

50.9 NOTES 50·24

50.1 INTRODUCTION.[1] Most security personnel have commiserated with colleagues about the difficulty of getting people to pay attention to security policies—to comply with what seems like good common sense. They shake their heads in disbelief as they recount tales of employees who hold secured doors open for their workmates—or for total strangers, thereby rendering million-dollar card-access systems useless. In large organizations, upper managers who decline to wear their identification badges discover that soon no one else will either. In trying to implement security policies, practitioners sometimes feel that they are involved in turf wars and personal vendettas rather than rational discourse.

These problems reflect the social nature of human beings; however, they also reflect the fact that although people involved in information systems security and network management have a wide variety of backgrounds, many lack training in social or organizational psychology.

Security policies and procedures affect not only what people do, but also how they see themselves, their colleagues, and their world. Despite these psychosocial issues, security personnel pay little or no attention to what is known about social psychology. The established principles of human social behavior have much to teach in any attempts to improve corporate and institutional information assurance (IA).

IA specialists concur that security depends on people more than on technology. Another commonplace is that employees are a far greater threat to IA than outsiders (see Chapter 13 in this *Handbook*).

It follows from these observations that improving security necessarily involves changing beliefs, attitudes, and behavior, both of individuals and of groups. Social psychology can help us understand how best to work with human predilections and predispositions to achieve our goals of improving security:

- Research on social cognition looks at how people form impressions about reality. Knowing these principles, we can better teach our colleagues and clients about effective security.
- Work on attitude formation and beliefs helps to present information effectively and so convince employees and others to cooperate in improving security.
- Scientists studying persuasion and attitude change have learned how best to change people's minds about unpopular views, such as those regarding the security community.
- Studies of factors enhancing prosocial behavior provide insights on how to foster an environment where corporate information is willingly protected.
- Knowledge of the phenomena underlying conformity, compliance, and obedience can help to enhance security by encouraging compliance and by protecting staff against social pressure to breach security.
- Group psychology research provides warnings about group pathology and about hints for working better with groups in establishing and maintaining IA in the face of ingrained resistance.

This chapter reviews well-established principles of social psychology that help security and network management personnel implement security policies more effectively. Any recent introductory social psychology college textbook will provide ample references to the research underpinning the principles applied here to security policy implementation.[2]

50.2 RATIONALITY IS NOT ENOUGH. IA policies sometimes evoke strong emotions. People can get very angry about what they perceive as interference with their way of getting their work done. From the perspective of the traditional information security professional, information security is still perceived as a technical problem, but as recent research reveals, it is more of a management problem, and the prevalent security culture offers insight into how management handles this problem.

To put the discussion in context, here is a definition of rationality from a respected academic, Jonathan Baron: Rationality is "the kind of thinking we would all want to do, if we were aware of our own best interests, in order to achieve our goals."[3]

Applying this definition to the concept of security, it seems to imply that rational thought would direct us to what appears to be the best compromise between what we perceive as our security needs and what appears to be the most convenient process.

50.2.1 Schema. Psychologists use the word *schema* to summarize the complex picture of reality on which we base our judgments. The schema is what social psychologists call the way people make sense of their social interactions. IA practitioners must often change their colleagues' schemata.

Schemata are self-consistent views of reality. They help us pay attention to what we expect to be important and to ignore irrelevant data. They also help us organize our behavior. For example, our schema for relations at the office includes polite greetings, civil discussions, written communications, and businesslike clothes. The schema excludes obscene shrieks, abusive verbal attacks, spray-painted graffiti, and colleagues dressed in swimsuits. It is the schema that lets people know what is appropriate or inappropriate in a given situation.

Unfortunately, security policies and procedures conflict with most people's schemata. Office workers' schemata includes sharing office supplies ("Lend me your stapler, please?"), trusting their team members to share information ("Take a look at these figures, Sally"), and letting their papers stay openly visible when they leave their desks.

Sharing user IDs, showing sensitive information to someone who lacks the appropriate clearance, and leaving workstations logged on without protection are gross breaches of a different schema—that of the IA specialist. Think about access controls: Normal politeness dictates that when a colleague approaches the door we have just opened, we hold the door open for the person; when we see a visitor, we smile politely—after all, it might be a customer. In contrast, access-control policies require that we refuse to let even well-liked colleagues piggyback their way through an access-card system; security policies insist that unbadged strangers be challenged or reported to security personnel. Common sense tells us that when the chief executive officer (CEO) of the company wants something, we do not oppose it; yet good IA dictates that we train computer room operators to forbid entry to anyone without documented authorization—including the CEO.

Sometimes people subvert IA by systematically getting around the rules because their normal social schema supersedes the security schema. It is not uncommon for naïve staff to give keys or the door lock combination for access into secured areas to regularly schedule outside delivery and maintenance persons. Such delivery people are rarely subjected to security checks, and yet their potential for intentional or inadvertent damage is great; nonetheless, the naïve staff members are acting without authorization, subverting normal security controls, and entrusting the safety and security of corporate resources to relative unknowns: Why? Are they deliberately violating security policy with evil intent? Of course not: The employees are simply acting in a friendly fashion and extending trust that might be appropriate in other circumstances—but they are using the wrong schema for a high-security corporate environment. In contrast, an IA specialist's schema in the same circumstances includes all the potentially untrustworthy friends of those outsiders. Until the security administration group alters the employees' perception of the appropriateness of the security regulations, the conflict between different schemata will continue to cause security violations and fuel resentment on all sides.

Indeed, a common response to attempts at enforcing existing policies and procedures, or to new security rules, is a charge of paranoia leveled against security personnel. Other accusations include authoritarian behavior and undue interference with job functions. These responses usually indicate a conflict between accepted norms of behavior and the need to change behavior to conform to security principles. They imply that social and cultural needs and behavior were not accounted for in the design of the principles. They also indicate a need for security personnel to understand that basic social graces and accepted norms of social activity conflict with the basic needs of security—and that the employees violating security rules are not inherently bad people.

Some redesign might allow for better security by recognizing that security rules can violate accepted social norms. For example, in the case of security-locked doors, one could allow for relatively free access to common areas (thus allowing people to hold the door for colleagues following them) while forbidding such actions in secured locations. A posted explanation of both the reasons for, and consequences of, the policy explicitly addressing the difference between the needs of normal politeness and the needs of high security could lead to better adherence to stated policy. The text might read something like this:

> This is a high-security area. Preventing anyone from entering without swiping their owned access card is not rude: it's common sense. Entering secured areas without using each employee's access card could put employees who did not use their access cards to enter at risk in an emergency if security staff didn't know that they were still in the secured areas. You are welcome to exercise normal politeness by holding the doors open for your colleagues and badge-wearing visitors in the nonsecured areas.

If we persist in assuming that we can influence our colleagues to change their perception of IA solely by informing, cajoling, nagging, or browbeating them, we will continue to fail. IA must be integrated into the corporate culture by changing our colleagues' schemata, a process that needs to use all of the techniques that social psychology can teach us.

A simple measure of this reality is to be found in the persistent avoidance practiced by many U.S. Government agencies in applying the information security program requirements of the Federal Information Security Management Act of 2002 (FISMA). It took almost five years and a series of undesirable security incidents before the various agencies woke up to the reality of the need to implement appropriate protection.[4] Another illustration of the reluctance to implement security policies is the Veterans Affairs debacle in the late 2000s involving loss of control over personally identifiable information on unencrypted disk drives.[5]

50.2.2 Theories of Personality. One of the most pervasive obstacles to cooperation in organizations is interpersonal conflict. Many conflicts are rooted in differences of *personality style*. For example, one widely used set of categories for describing people's personalities uses this schema:

- Extroversion
  - High: active, assertive, energetic, outgoing, talkative
  - Low: quiet, reserved, shy, silent, withdrawn
- Agreeableness
  - High: affectionate, appreciative, kind, soft-hearted, sympathetic
  - Low: cold, fault-finding, hard-hearted, quarrelsome, unfriendly

- Conscientiousness
  - High: efficient, organized, planful, responsible, thorough
  - Low: careless, disorderly, frivolous, irresponsible, slipshod
- Emotional stability
  - High: calm, contented, stable, unemotional
  - Low: anxious, moody, nervous, tense, worrying
- Openness or culturedness
  - High: imaginative, insightful, intelligent, original, wide interests
  - Low: commonplace, shallow, simple, narrow interests, unintelligent

The adjectives used in this summary are positive for the high side of each trait and negative for the low side. However, the assumption that different personality types are easily characterized as superior and inferior seriously interferes with respectful communications among colleagues. For example, people with low characteristics might view the preceding summary in this way:

- Extroversion
  - High: nervous, aggressive, excitable, pushy, chattering
  - Low: dignified, respectful, unassuming, attentive, self-sufficient
- Agreeableness
  - High: clinging, gushy, soft-headed, knee-jerk reactive, uncritical
  - Low: stately, analytical, rational, principled, reserved
- Conscientiousness
  - High: obsessive, compulsive, unspontaneous, pompous, slavish
  - Low: free, spontaneous, creative, fun, youthful, having perspective
- Emotional stability
  - High: frozen, ambitionless, boring, dead
  - Low: vibrant, romantic, alive, strong, sensible
- Openness or culturedness
  - High: flaky, theoretical, complicated, off-the-wall, dilettante
  - Low: earthy, smart, grounded, focused, practical

In discussing corporate culture change, leaders must be on guard to defuse conflicts based on the misperception that one particular response or view of an issue is necessarily good and another necessarily bad. The conflict may be rooted in personality styles rather than in problems of understanding. If the security working group proposes that all employees must challenge anyone in the secured areas who is not wearing a badge, some people—those who have low extroversion, for example—may have a great deal of difficulty with the concept that they should tell anyone else what to do, especially a manager of a higher rank than their own. Arguing only over the reasons why such a policy would be *useful* would sidestep the fundamental problem: that the required behavior is in direct conflict with possibly lifelong and firmly held views on appropriate behavior.

Security personnel must remember that failure to comply with policy is not necessarily the result of a bad attitude. When it becomes obvious that conflicts are rooted in

personality, security personnel will have to try to arrive at a useful compromise. Instead of requiring that everyone confront the unbadged individual personally, the security policies could include a proviso allowing for individuals to choose simply to inform security personnel immediately.

Role-playing exercises sometimes can defuse a problem in accepting security policies by desensitizing resistant personnel. Going through the motions of what they fear or dislike sometimes can help them come to realize that the proposed change in behavior is not as bad as they originally thought. Returning to the example of confronting violations of security, many people have difficulty imagining that they could tell a superior in the management hierarchy not to *piggyback*. This term, like *hitchhiking* and *tailgating*, describes entering through a secured door that has been opened by someone else using a valid access code or token. Going through exercises in which each person pretends in turn to be the upper manager and then the challenger helps to break down resistance to this particular security policy. Trainers can emphasize that the challenge must be seen using a different schema from the normal situation:

- It is socially acceptable to apply security policies within the organization.
- Higher-status employees can encourage lower-status employees by articulating their support for the policy and showing that they are not offended by the request.
- Participants of all hierarchical levels can introspect to see that they themselves do not feel offended when they are politely asked to use their badge during the role-playing exercise.

In general, leaders of the security team responsible for implementing security policies should be on the lookout for conflicts of style that interfere with the central task of making the enterprise more secure. If an individual likes short, direct instructions without chitchat about nonessentials, the security team member should adapt and stick to essentials; if an individual is known to like getting to know a stranger and wants to spend a few minutes learning about family background, it should not be opposed. Communicating ideas in a way that is likely to be acceptable is more important than imposing one's own interpersonal style preferences on others.

Above all, security personnel—and management in general—ought to be doing a great deal more *listening* and a great deal less *commanding.*

Some important psychological issues for security leaders to consider include:

- Digital security is extremely complicated, and the explanations can be very technical—both attributes that are unfavorable to fostering management attention.
- Most security incidents stem from insiders rather than from outsiders, so prevention requires consistent nagging—not something management or anyone else normally regards favorably.
- Success in digital security (in fact, all of security) is best shown by having *nothing* happen, which is a tough thing to measure and tougher to sell. So the personal payoff for a well-executed security strategy is often little to nothing—and no management executive wants to put up *nothing* as a true measure of success.

The practical implications of these observations include:

- Discussions of IA should be down-to-earth and practical whenever possible.
- Awareness can be achieved by more positive means than nagging.

- One can create metrics of success by using security games and competitions that are fun as well as informative and effective at maintaining security awareness.

Here is an abstract of an article on the general manager's contribution to security:

> Few senior executives pay much attention to computer security. They either hand off responsibility to their technical people or bring in consultants. But given the stakes involved, an arm's-length approach is extremely unwise. According to industry estimates, security breaches affect 90% of all businesses every year and cost some $17 billion. Fortunately … senior executives don't need to learn about the more arcane aspects of their company's IT systems to take a hands-on approach. Instead, they should focus on the familiar task of managing risk. Their role should be to assess the business value of their information assets, determine the likelihood that those assets will be compromised, and then tailor a set of risk abatement processes to their company's particular vulnerabilities. *This approach, which views computer security as an operational rather than a technical challenge, is akin to a classic quality assurance program in that it attempts to avoid problems rather than fix them and involves all employees, not just IT staffers.* The goal is not to make computer systems completely secure—that's impossible—but to reduce the business risk to an acceptable level…[6]

We have italicized the key sentence in the quote to emphasize the critical role of changing corporate culture in successful security management.[7]

50.2.3 Explanations of Behavior. In practice, trying to change corporate culture can be a frustrating and long-drawn-out project. One aspect of this process that security group leaders should monitor closely is the interpretation of employee behavior (called *attribution* theory in the social psychology literature) by members of the security team. In general, people can be viewed as interpreting (i.e., explaining) other people's behavior according to two independent dimensions: *internal or external* and *stable or unstable*. Here are some explanations of why Betty has failed to log off her session for the fourth time this week before leaving the office:

- **Internal, stable.** "That's just the way she is—she never pays attention to these rules."
- **Internal, unstable.** "She's been under strain lately because her child is sick—that's why she's forgotten."
- **External, stable.** "The system doesn't respond properly to the logoff command."
- **External, unstable.** "This week, the system has not been responding properly to the logoff command."

This simple four-way classification is useful for leaders in understanding and avoiding classic errors of attribution. Such attribution errors can cause conflicts between the security staff and other employees, or even among employees with different degrees of compliance to policy.

50.2.4 Errors of Attribution. Some well-established misinterpretations of others' behavior can interfere with the acceptance of security policies. Such errors interfere with the ability of security personnel to communicate the value of security policies. Security group leaders should sensitize their staff to the consequences of these errors.

50.2.4.1 Fundamental Attribution Error. The most important error people use when explaining other people's behavior is to assume that a person's actions are stable, internal features; a typical example of this error is the naïve belief that an actor's personality is essentially what that person portrays in performance. Anyone who has ever experienced surprise at the demeanor and speech of a favorite actor who is being interviewed has committed the *fundamental attribution error*. Some actors who play bad characters in fictional situations have even been verbally and physically assaulted by viewers who cannot resist the fundamental attribution error, and who genuinely believe that the actors are as bad as the nasty people they portray. The abstract of a conference presentation in 2008 summarized some significant findings illustrating this point:

> Two studies attempted to document the occurrence of the psychological phenomenon known as the fundamental attribution error (FAE) in the audiovisual medium. The FAE refers to the human tendency to attribute people's behavior to internal attributes more than external factors. In Study 1, we demonstrated that in the audiovisual medium, viewers tend to attribute an actor's behavior in television dramas to the actor's personality, ignoring the existence of a script dictating the actor's behavior. Study 2 replicated this finding, and also demonstrated that the tendency to make the FAE is related to the degree to which the person reports being transported into the narrative of the TV drama. Furthermore, we showed that the tendency to attribute character traits to the actor is not diminished following exposure to the same actor playing two opposing roles. The last scene viewed was found to determine the evaluation of the actor's characteristics.[8]

In security work, being on guard against the fundamental attribution error helps to smooth relations with other employees. For example, if a security group member sees an employee, Jill, who is not wearing her badge, it is easy to assume that she never wears her badge and is refusing to wear it because of a character flaw. The security officer may act according to these assumptions by being harsh or unfriendly in correcting Jill's behavior. The harshness generates resentment, and Jill may come to associate security with unpleasant people, thus reducing the likelihood that she will comply with policy or encourage others to do so.

In fact, however, most people's behavior is far less stable and internal than unstable and externally based. For example, if the security officer simply smiled and pointed gently to the lack of a badge instead of jumping to conclusions, he might discover that Jill's lack of a badge today was due simply to her having taken her jacket off just before an urgent call from the vice president, interrupting her normal procedure of moving the badge from jacket to shirt pocket. Thus, her lack of a badge would not be stable behavior at all—it would be a temporary aberration of no longlasting significance. A solution would be to get used to clipping the badge to her trousers or her skirt instead of her jacket or her blouse. Similarly, just by asking nicely, the security officer might learn that Jill normally does wear her badge, but today her four-year-old son took it off her jacket to play with it, without his mother's noticing the change. In this example, Jill's behavior is externally based and has nothing to do with character. The kindly interactions between Jill and the security guard *increase* the sense of social relation and make it more likely that she will remember the incident positively and comply with the security policy in the future.

In summary, by being aware of the fundamental attribution error, security personnel can be trained to adopt a less judgmental, or quick-draw, mentality that can alienate other employees and damage security programs.

50.2.4.2 Actor-Observer Effect. The *actor-observer effect* consists of interpreting one's own behavior as appropriate unstable, externally motivated responses to environmental conditions, whereas other people's behavior is viewed in the light of the fundamental attribution error as stable, internally motivated expressions of character traits. Becoming aware of this tendency helps security personnel resist the fundamental attribution error.

50.2.4.3 Self-Serving Bias. The counterpart of the actor-observer effect is the *self-serving bias*, which fools people into believing that their own behavior is due to stable, internal aspects of their character. Security officers who are unaware of this dangerous error may come to feel that they are in some sense superior to other people who do not know as much about security as they do or who do not comply as fully as they do with security policy. The officers may have failed to integrate the fact that hours of training and coaching by their security group leaders are at least as responsible for their own knowledge of, and compliance with, security policies as any innate superiority.

By bringing this kind of erroneous thinking to light during training and supervision of security staff, managers can help reduce the conflicts that naturally result from an air of assumed superiority.

50.2.4.4 Salience and Prejudice. When people are asked to guess which person in a group is the most influential (or least influential) person, social psychologists find that whichever person stands out the most, for whatever reason, is more often attributed with the special properties in question. Such effects apply to any characteristic that the psychologists ask about: most (or least) intelligent, aggressive, sympathetic, and so on. This phenomenon is known as the *salience effect.*

An application of the salience effect might occur if security officers see a group of employees who are violating security policies. A natural and counterproductive tendency is to leap to the conclusion that the tallest or shortest, the thinnest or fattest, the whitest or blackest person in the group must be to blame. This error can result in unfair treatment of perfectly innocent people.

This problem of misinterpreting salience is exacerbated by prejudice; for example, imagine there were an identifiable group called the *Ogunians* (as far as we can determine, there is no such group) who traditionally wear, say, a seven-sided symbol of their identity. If an anti-Ogunian security officer sees a noncompliant group where one of the members is wearing the characteristic heptagon of Ogun, it may be hard for the officer to resist blaming the noncompliance on the Ogunian even if, in fact, the Ogunian was waiting to use a valid access card in full compliance with security policy.

Worse, people can be so strongly influenced by expectation—part of their schema—that they actually misperceive a situation altogether. For example, in some classic experiments studying prejudice in the 1950s, psychologists showed subjects a drawing of two people, one light-colored and the other dark-colored, standing in a tramway car. One was holding a knife. When questioned about the image afterward, white subjects consistently reported that the black figure had been holding the knife, but actually it was the white figure in the drawing who had the knife.

Thus, even observation itself can be twisted by prejudice and expectations; for example, if the anti-Ogunian security officer sees a group of people passing through an open doorway into a secured area without using their badges, the officer may incorrectly

report that it was the fault of an Ogunian even if there was no Ogunian in the group. Such a mistaken report would not only infuriate innocent Ogunians and possibly cause general Ogunian resentment or hostility toward security efforts in general, but it also could mislead the security group itself into trying to correct the behavior of the wrong person or people.

Similarly, any minority—whether in terms of gender, gender orientation, religion, race, or disability—can be the focus of a prejudiced security officer's blame when a group disobeys policy. Security leaders should make their staff aware of the danger of applying this erroneous method of explaining group behavior. In many organizations, such discrimination is a violation of corporate policy and may even be illegal. In any case, prejudice is not constructive and must be monitored and overcome.

50.2.5 Intercultural Differences. Many countries in the world are experiencing changes in their population due to immigration. Especially in areas where people have heretofore been largely homogeneous, cultural, religious, and racial diversity can lead to interpersonal and intergroup conflicts. Such conflicts may be based in part on prejudice, but they also may be the result of differing values and assumptions.

This definition of culture helps define the discussion:

> **Culture as Mental Programming**
> Every person carries within him- or herself patterns of thinking, feeling, and potential acting that were learned throughout their lifetime. Much of it has been acquired in early childhood, because at that time a person is most susceptible to learning and assimilating. As soon as certain patterns of thinking, feeling, and acting have established themselves within a person's mind, he or she must unlearn these before being able to learn something different, and unlearning is more difficult than learning for the first time.
>
> Using the analogy of the way computers are programmed, this book will call such patterns of thinking, feeling and acting *mental programs,* or … *software of the mind*…
>
> A customary term for such mental software is *culture.*[9]

Security personnel engaged in the process of corporate culture change should be sensitive to the possibility that people with different real-world cultural backgrounds can respond differently to proposed security policies. For example, in 2001 the fundamentalist extremists of the Taliban in Afghanistan decreed that non-Muslim people would have to wear badges in public.[10] One can imagine that a Hindu Afghan refugee in the United States who is told to wear a badge for security reasons might have an unexpectedly emotional response to the order. Before pressuring (or becoming hostile to) anyone who seems to be resisting a policy, it is valuable to inquire about the person's beliefs and attitudes and to explain the foundation for the policies in question. Especially where there are intercultural differences, such inquiry and discussion can forestall difficulties and dissension and assuage unexpected, culturally rooted anxiety.

Security professionals need to be acutely aware of the cultural differences of individuals in their target audiences. For example, in some cultures (mostly the new world and western), reality is directly related to facts and verifiable calculations. Other cultures may put a stronger emphasis on personal feelings, intuition, and culturally ingrained beliefs into their comprehension of reality. Hence different people may not share the same understanding of reality or implement policies by the same principles unless they make their assumptions known and discuss them with the intention of coming to agreement.

When considering culture, the changing dynamics of modern society also need to be accounted for. Any modern city or leading public institution today is a complex

combination of people from different cultural and social backgrounds. Working together for large parts of their active days together as a group leads to a melting pot situation, with all the individual cultures and social leanings being added into the mix and emerging as a (mostly) different culture. The individuals then take this new culture back home, where it mixes in with their family and friends' contributions, resulting in often more modifications and variation. The organization itself is constantly changing, as people join and leave. All of this leads to culture being a dynamic phenomenon, albeit not as rapidly changing as the technology front. Hence, any security policy or framework needs to be able to account for and accommodate this changing paradigm if it has to be successful in securing the enterprise.

50.2.6 Framing Reality. How can we make the corporate culture more supportive of IA?

Schemata influence what we perceive. For example, an employee refuses to take vacations, works late every night, is never late, and is never sick. A model employee? Perhaps, in one schema. From the security point of view, the employee's behavior is suspect. There have been cases where such people have been embezzlers unable to leave their employment: Even a day away might result in discovery of their crimes. Saint or sinner? Our expectations determine what we see.[11]

To change the schema so that people take IA seriously, we should provide participants in training and security awareness with real-life examples of computer crime and security breaches, so that security policies make obvious sense rather than seeming to be arbitrary.

Schemata influence what we remember. When information inconsistent with our preconceptions is mixed with details that fit our existing schemata, we selectively retain what fits and discard what conflicts. When we have been fed a diet of movies and television shows illustrating the premise that information is most at risk from brilliant hackers, why should we remember the truth: that carelessness and incompetence by authorized users of information systems cause far more harm than evil intentions and outsiders ever do?

Instructors should emphasize the practical side of IA by showing how policies protect all employees against false accusations, prevent damage to the organization's reputation and profits, and even play a role in national security. This is especially true where business touches the technical infrastructure on which we all depend.

Most important of all, teaching others about IA cannot be an occasional and haphazard affair. Before attempting to implement policies and procedures (aside from emergency measures that are needed at once), we should ensure that we build up a consistent view of IA among our colleagues. In light of the complexity of social cognition, our usual attempts to implement security policies and procedures seem pathetically inept. A couple of hours of lectures followed by a video, a yearly ritual of signing a security policy that seems to have been written by Martians—these are not methods that will improve security. These efforts merely pay lip service to the idea of security.

According to research on counterintuitive information, people's judgment is influenced by the manner in which information is presented. For example, even information contrary to established schemata can be assimilated if people have enough time to integrate the new knowledge into their worldviews. It follows that nonemergency security policies should be introduced over a long time, not rushed into place.

An effective IA program includes frequent reminders of security. To change the corporate culture, practitioners should use methods such as a security corner in the corporate publication, security bulletins detailing the latest computer crime or security

breach that has hit the news, contests for identifying the problems in realistic scenarios, and write-in columns to handle questions about policies. IA has to become part of the framework of reality, not just an imposition from management.

In every security course or awareness program, instructors and facilitators should explicitly address the question of corporate culture, expectations, and social schemata. Do not rely solely on intellectual discourse when addressing a question of complex perceptions and feelings. Use simulations, videos, and role-playing exercises to bridge the gap between intellect and emotion.

Address the feelings and perceptions of all participants as they learn about the counterintuitive behaviors that improved security will demand. Encourage learners to think about how they might feel and respond in various situations that can arise during the transition to a more secure environment. For example, ask participants to imagine:

- Asking colleagues not to step through a secured entrance without passing through the access-control system with their own identity
- Telling their boss that they will not copy software without a license to do so
- Questioning a visitor or employee who is not wearing an identity badge[12]

50.2.7 Getting Your Security Policies Across. What are some ways to change our colleagues' schemata so that they become more receptive to IA policies?

- **Initial exposure.** Preliminary information may influence people's responses to information presented later. For example, merely exposing experimental subjects to words such as *reckless* or *adventurous* affects their judgment of risk-taking behavior in a later test.

  It follows that when preparing to increase employee awareness of security issues, presenting case studies is likely to have a beneficial effect on participants' readiness to examine security requirements.
- **Counterexamples**. Preexisting schemata can be challenged by several counterexamples, each of which challenges a component of the schema. For example, prejudice about an ethnic group is more likely to be changed by contact with several people, each of whom contradicts a different aspect of the prejudiced schema.

  It follows that security awareness programs should include many realistic examples of security requirements and breaches. In a counterexample, students in college IA courses have commented on the unrealistic scenario in a training video they were shown: a series of disastrous security breaches occurring in the same company. Based on the findings of cognitive social psychologists, the film would be more effective for training if the incidents had been dramatized as occurring in different companies.

  In practical terms, practitioners should stay current and update their materials. Many IA publications provide useful case studies that will help make awareness and training more effective.
- **Choice of wording.** Perceptions of risks and benefits are profoundly influenced by the wording in which situations and options are presented. For example, experimental subjects responded far more positively to reports of a drug with "50 percent success" than to the same drug described as having "50 percent failure."

  It follows that practitioners should choose their language carefully during security awareness campaigns. Instead of focusing on reducing failure rates (violations

of policy), we should emphasize improvements in our success rates. Unfortunately, some rates cannot be expressed in positive terms; for example, it is not easy to measure the success rate of security measures designed to foil attacks on systems.

Judgments are easily distorted by the tendency to rely on personal anecdotes, small samples, easily available information, and faulty interpretation of statistical information. Basically, we humans are not always rational processors of factual information. If security awareness programs rely strictly on presentation of factual information about risks and proposed policies and procedures, they are likely to run up against a stubborn refusal to act logically. Security program implementation must engage more than the rational mind. We must appeal to our colleagues' imagination and emotion as well. We must inspire a commitment to security rather than merely describing it.

50.2.8 Reward versus Punishment. When enforcing security policies, too many organizations focus entirely on punishing those who break the rules. However, everything we know about modifying behavior teaches us to use reward rather than punishment. Punishing people who do not comply with security rules often generates resentment and hostility that carry over into future interactions. Instead of seeing information assurance as a benefit to the organization and to their interests, victims of harsh treatment can resist even well-intentioned, sensible changes in security policies simply because of the emotional overlay associated with the embarrassment and frustration generated by criticism and penalties.

In addition to avoiding negativity and push-back, reward may simply work better than punishment at changing behavior. For example, a security officer from a large corporation experimented with reward and punishment in implementing security policies. Employees were supposed to log off their mainframe terminals when leaving the office, but compliance rates were only around 40 percent. In one department, the security officer used the usual techniques recommended in the literature and common among security professionals; for example, she put up nasty notes on terminals that were not logged off, changed the passwords on delinquent accounts, and humiliated violators by forcing them to report to their bosses for authorization to obtain a new password. However, in a different department, she simply left a Hershey's Chocolate Kiss on the keyboard of every terminal whose user had indeed logged off before leaving. After one month of these two strategies, compliance rates in the department subject to punishment had climbed to around 60 percent. Compliance in the department getting chocolates had reached around 80 percent—and their feelings toward security were much more favorable than the norm.

This case illustrates some of the benefits of reward:

- Compliance rates were significantly higher than in the group subjected to punishment.
- Attitudes (see Section 50.3) are more likely to be positive.
- Costs can be lower because small rewards delivered en masse may be much cheaper and quicker to apply than administrative procedures applied one by one through management intervention.

50.3 BELIEFS AND ATTITUDES. Psychologists distinguish between beliefs and attitudes. A *belief* refers to cognitive information that need not have an emotional component. An *attitude* refers to an evaluation or emotional response. Thus, a person

may believe correctly that copying a large number of proprietary software packages without authorization is a felony while nonetheless having the attitude that it does not matter to him. A rational employee may *believe* that malware can be downloaded onto company computers through unauthorized software available on unvetted Websites (and thus answer a questionnaire evaluating security awareness correctly) yet have the *attitude* that the risk is negligible—and cheerfully go on downloading unauthorized software at work.

50.3.1 Beliefs. Beliefs can change when contradictory information is presented, but some research suggests that it can take up to a week before significant shifts are measurable. Other studies suggest that when people hold contradictory beliefs, providing an opportunity to articulate and evaluate those beliefs may lead to changes that reduce inconsistency.

These findings imply that corporate security must explore the current structure of beliefs among employees and managers. Questionnaires, focus groups, and interviews may not only help the security practitioner, they actually may help move the corporate culture in the right direction. The *Hawthorne Effect* is the name given to improvements in measured behavior resulting simply from employee responses to being studied; done correctly, honestly, and nonpunitively, inquiring into employee beliefs and attitudes may communicate a genuine interest by management in improving security policy and practice with input from everyone involved.

50.3.2 Attitudes. An attitude, in the classical definition, is a learned evaluative response, directed at specific objects, which is relatively enduring and influences behavior in a generally motivating way. The advertising industry spends over $50 billion yearly to influence public attitudes in the hope that these attitudes will lead to changes in spending habits—that is, in behavior.

Research on classical conditioning suggests that attitudes can be learned even through simple word association. If we wish to move our colleagues toward a more negative view of computer criminals, it is important not to portray computer crime using positive images and words. Movies that show criminal hackers as pleasant, smart, physically attractive, and likable people may do harm by minimizing the seriousness of industrial espionage and cybervandalism. When teaching security, we should avoid praising the criminals we describe in case studies.

Studies of how attitudes are developed consistently show that rewards and punishments are important motivators of behavior. Studies show that even apparently minor encouragement can influence attitudes. A supervisor or instructor should praise any comments that are critical of computer crime or that support the established security policies. Employees who dismiss security concerns, or who flout the regulations, should be challenged on their attitudes, not ignored. Such challenges are best carried out in private to avoid causing embarrassment to the skeptics and possibly generating resistance due to pride or a sense of machismo.

50.3.3 Changing Attitudes toward Security. Persuasion—changing someone's attitudes—has been described in terms of communications. The four areas of research include:

1. **Communicator variables.** Who is trying to persuade?
2. **Message variables.** What is being presented?

3. **Channel variables.** By what means is the attempt taking place?
4. **Audience variables.** At whom is the persuasion aimed?

50.3.3.1 Communicator Variables. Attractiveness, credibility, and social status have strong effects immediately after the speaker or writer has communicated with the target audience; however, over a period of weeks to a month, the effects decline until the predominant issue is message content. We can use this phenomenon by identifying the senior executives most likely to succeed in setting a positive tone for subsequent security training. We should look for respected, likable people who understand the issues and sincerely believe in the policies they are advocating.

One personality style in particular can threaten the success of security policies: the *authoritarian personality.* A body of research suggests that some people, often those raised by punitive parents highly concerned with social status, become rigidly devoted to conventional beliefs, submit to authority, exercise authority harshly themselves, and are hostile to groups they perceive as unpopular. An authoritarian person might make a terrible security officer. Such an officer might derive more satisfaction from ordering people around and punishing them than from long-term success in implementing security policies.

50.3.3.2 Message Variables. Fear can work to change attitudes only if judiciously applied. Excessive emphasis on the terrible results of poor security is likely to backfire, with participants in the awareness program rejecting the message altogether. Frightening consequences should be coupled immediately with effective and achievable security measures.

Some studies suggest that presenting a balanced argument helps convince those who initially disagree with a proposal. Presenting objections to a proposal and offering counterarguments is more effective than one-sided diatribes. Popular training videos from the Software & Information Industry Association use this technique: they show people such as "college students, college faculty and publishers of all types of media discuss[ing] the legal and ethical implications of copying other people's works" and fairly present the arguments of copyright violators before rebutting them.[13]

Modest repetition of a message can help generate a more positive response. Thus, security awareness programs that include imaginative posters, mugs, special newsletters, audio and videotapes, and lectures are more likely to build and sustain support for security than occasional intense sessions of indoctrination. The use of multiple communications channels (discussed in the next section) also increases the effectiveness of the message.

50.3.3.3 Channel Variables. The channel through which we communicate has a strong effect on attitudes and on the importance of superficial attributes of the communicator. In modern organizations, most people assume that a meeting is the ideal way to communicate new information. However, the most effective medium for convincing someone to pay attention to any topic is face-to-face persuasion. Security training should include more than tapes and books; a charismatic teacher or leader can help generate enthusiasm for—or at least reduce resistance to—better security.

In addition, security educators should not introduce new ideas to decision makers in a meeting. There is too much danger of confounding responses to policy with nonpolicy matters rooted in relationships among the participants. It is not uncommon for one executive to oppose a new policy simply because another has supported it. A good way

to introduce security policies is to have individual meetings with one executive at a time in order to explain the issues and proposals and to ask for support.

Psychologists testing cognitive response theory have studied many subtle aspects of persuasion. Experiments have shown that rhetorical questions, such as "Are we to accept invasions of our computer systems?" are effective when the arguments are solid but counterproductive when arguments are weak. Security officers should not ask rhetorical questions unless they are certain that almost everybody will inevitably have the same answer—the one the security officers are looking for.

Consideration of facts and logical arguments, as the central route to persuasion, has been found to lead to more lasting attitudes and attitude changes than the peripheral influences from logically unrelated factors, such as physical attractiveness of a speaker.

50.3.3.4 Audience Variables. As mentioned, questionnaires and interviews may help cement a favorable change in attitude by leading to commitment. Once employees have publicly avowed support for better security, some will begin to change their perception of themselves. Specific employees should be encouraged to take on various areas of public responsibility for IA within their work group. These roles should periodically be rotated among the employees to give everyone the experience of public commitment to improved security.

To keep up interest in security, regular meetings of enthusiasts to discuss recent security news can keep the subject fresh and interesting. New cases can help security officers explain policies with up-to-date references that will interest their fellow employees and motivate managers to pay attention to security policies.

50.4 ENCOURAGING INITIATIVE. The ideal situation would be for everyone actually to help enforce security policies. Actually, however, some people are cooperative and helpful whereas others—or even the same people in different circumstances—are reluctant and suspicious about new policies. What can we do to increase cooperation and reduce rejection?

50.4.1 Prosocial Behavior. Studies of people who have come to the aid of others can help to encourage everyone in an organization to do the right thing. Some people intervene to stop crimes; others ignore crimes or watch passively. Social psychologists have devised a schema that describes the steps leading to prosocial behavior:

1. People have to notice the emergency or the crime before they can act. Thus, security training has to include information on how to tell that someone may be engaging in computer crime.
2. The situation has to be defined as an emergency—something requiring action. Security training that provides facts about the effects of computer crime on society and solid information about the need for security within the organization can help employees recognize security violations as emergencies.
3. Everyone must take responsibility for acting, but the larger the number of people in a group confronted with an emergency, the slower the average response time. Larger groups seem to lead to a *diffusion of responsibility*; each person feels that someone else is more responsible for dealing with the emergency. Another possible factor is uncertainty about the social climate; people fear appearing foolish or overly emotional in the eyes of those present. To overcome this effect, a corporate

culture must be established that rewards responsible individual behavior, such as reporting security violations.

4. Once responsibility for solving a problem has been accepted, appropriate decisions and actions must be taken. Clearly written security policies and procedures will make it more likely that employees act to improve security. In contrast, contradictory policies, poorly documented procedures, and inconsistent support from management will interfere with the decision to act.

Another analysis proposes that people implicitly analyze costs of helping and of not helping when deciding whether to act prosocially. The combination of factors most conducive to prosociality is low cost for helping and high cost for not helping.

Security procedures should make it easy to act in accordance with security policy. There should be a hotline for reporting security violations, and anonymity should be respected if desired. Psychological counseling and follow-up should be available if people feel upset about their involvement. Conversely, failing to act responsibly should be a serious matter; personnel policies should document clear and meaningful sanctions for failing to act when a security violation is observed. Penalties would include critical remarks in employment reviews and, where appropriate, even dismissal.

One method that does *not* work to increase prosocial behavior is *exhortation*; merely lecturing people in the abstract about what they ought to do has little or no positive effect.

Significantly, the general level of stress and pressure to focus on difficult tasks with seemingly impossible deadlines can greatly reduce the likelihood that people will act on their moral and ethical principles. Security is likely to flourish in an environment that provides sufficient time and support for employees to work professionally. Offices where everyone responds to a continuing series of apparent emergencies will not be likely to pay attention to security violations.

Some findings from research confirm common sense. For example, guilt motivates many people to act more prosocially. This effect works best when people are forced to assume responsibility. Thus, enforcing standards of security using reprimands and sanctions can indeed increase the likelihood that employees subsequently will act more cooperatively; however, as suggested earlier, punishment should not replace reward.

In addition, mood affects susceptibility to prosocial pressures. Bad moods make prosocial behavior less likely, whereas good moods increase prosociality. A working environment in which employees are respected is more conducive to good security than one that devalues and abuses them.

Even cursory acquaintance with other people makes it more likely that we will help them; it thus makes sense for security supervisors to get to know the staff from whom they need support. Encouraging social activities in an office (e.g., lunchtime discussion groups, occasional parties, and charitable projects) enhances interpersonal relationships and can improve the climate for effective security training. *Management by walking around* is an excellent practice at many levels, including fostering at least the first stage of interpersonal relationship among coworkers.[14]

50.4.2 Conformity, Compliance, and Obedience. Some people react negatively to the words *conformity*, *compliance*, and *obedience*, but ignoring social phenomena will not help security practitioners to attain their goals. Despite the unpopularity of this subject area, it is valuable to understand how people can work together in reinforcing security policies. The next sections look at how to increase conformity

with a culture of cooperation for increased security, compliance with rational security rules, and a bias toward obedience to IA authorities in security matters.

50.4.2.1 Social Pressure and Behavior Change. Turning a group into a community provides a framework within which social pressures can operate to improve an organization's IA. Most people respond to the opinions of others by shifting their own opinions, sometimes unconsciously, toward what statisticians call the *mode*—the most popular opinion. Security programs must aim to shift the normative values, the sense of what one should do, toward protecting confidentiality, possession or control, integrity, authenticity, availability, and utility of data.

According to an informal survey conducted by Mani Akella, a coauthor of this Chapter, at three leading financial firms on Wall Street, these inferences resulted from a test group of 80 respondents:

- Older employees prefer to model their reactions based on common group preferences, even if some of the reactions go against their own gut feeling. The rationale here seems to be that the group provides anonymity and even insulates them from management reaction. Younger employees, however, tend to buck the group trend when they disagree with proposed concepts.
- Leadership has a large role to play—and the group modifies its reactions very quickly to adapt to leadership changes. If the leader likes to follow a specific path and not ask questions, the entire group tends to let issues lie and not disturb the even tenor of the organization for fear of disturbing the leader, even at the cost of risking serious potential security lapses (see Section 50.5.4 on groupthink). If the leader fosters a dynamic, open, and collaborative environment with a measured adaptability to evolving threats, however, the group enlivens itself with innovation and puts out additional effort to stay abreast (or even ahead) of the current threat landscape.
- When the leaders challenge the individuals to greater achievement without threats of punitive reaction, the *group* reacts with positive response. Leadership can create a security environment that can exceed the enterprise's security expectations by encouraging the individual to increase productivity and to reward oneself with greater job satisfaction. Security, like most other organizational management efforts, is all about people. Responsible, satisfied, and aware personnel naturally lead to better overall security for the organization.

50.4.2.2 Changing Expectations. As has been evident in public campaigns aimed at eliminating drunken driving, it *is* possible to shift the mode. In the United States in the mid-twentieth century, many people believed that driving while intoxicated was amusing; today, a drunken driver is a social pariah. High school children used to kill themselves in large numbers on the nights of their high school proms; today, many children spontaneously are arranging for safe rides home. In much the same way, we must move toward making computer crime as distasteful as public drunkenness.

The trend toward similar behavior increases when people within the group like or admire each other. In addition, the social status of an individual within a group influences that individual's willingness to conform to group standards. High-status people (those liked by most people in the group) and low-status people (those disliked by the group) both tend to be more autonomous and less compliant than people liked by some and disliked by others. Therefore, security officers should pay special attention

to those outliers during instruction programs. Managers should monitor compliance more closely at both ends of the popularity range. If security practices are currently poor, and allies are needed to change the norm, working with the outliers to resist the majority's anti-security bias may be the most effective approach. The most popular people may be disastrous agents of rebellion if they do not sign on to the security program; paradoxically, the most unpopular people may be helpful if they can be persuaded to comply.

50.4.2.3 Norm of Reciprocity. According to social psychologists, the norm of reciprocity indicates that, in social relations, favors are usually returned. Even a small, unexpected, unsolicited, or even unwanted gift increases the likelihood that we will respond to requests. For example, members of various religious cults often hand out flowers or books at airports, knowing that the norm of reciprocity will increase the frequency and amount of donations from basically uninterested passersby.

A security awareness program that includes small gifts, such as an attractive mug labeled "SECURITY IS EVERYONE'S BUSINESS" or an inexpensive but useful booklet summarizing security policies, can help get people involved in security. The combination of such programs with rewards for compliance can be a powerful tool for improving security.

Combining a token of appreciation with direct personal contact starting with the statement "I need your help" followed by a frank exposition of the security situation can be positive at all levels. This approach works at multiple levels—establishing personal relations, building on the norm of reciprocity, and changing the schema.

50.4.2.4 Incremental Change. The foot-in-the-door technique suggests that a small initial request should be followed by an even larger second one. Political field workers, for example, know that they can start small by asking people to let them put candidate stickers in their window; then they ask to put a candidate's poster on their lawn; eventually they can ask for volunteer time or money. Every compliance with a request increases the likelihood that the person will agree to the next step in an escalating series. It is as if agreeing to one step helps to change the targets' sense of themselves. To reduce discomfort about their beliefs and their behavior (what psychologists call *cognitive dissonance*), people change their beliefs to conform with their behavior.

Employees can be asked personally to set a good example by blanking screens and locking terminals when leaving their desks. Later, once they have begun the process of redefining themselves ("I am a person who cares about computer security"), they can be asked for something more intense, such as participating in security training by asking others to blank their screens and lock their terminals—or rewarding those who do with the famous chocolate tidbit. By applying the same methods to various tasks, the corporate culture can change so that a majority of people feel personally committed to good security practices.

Some security specialists have proposed that we should not ask the audience to think. The reasoning is that each incremental policy step should not require the target audience to reason or explain behavior. Rather, focus on building conditioned reflexes to specific environmental and usage factors. The organization and the security team should be more assured of a common and predictable reaction to any security threat from each individual internal person. However, a countervailing view is that every behavior proposed to improve security *must* be grounded in an understandable schema. In other words, although one need not force the members of the audience to *articulate*

the rationale, the rules must make sense if they are to be integrated, remembered, and applied in the long term.

50.5 GROUP BEHAVIOR. Some groups of people are referred to as *teams*, while others are called *gangs*. Social psychological insights into group behavior can improve success rates for IA policies.

50.5.1 Social Arousal. Studies on the behavioral effects of being in groups produced contradictory results; sometimes people did better at their tasks when there were other people around, and sometimes they did worse. Eventually, psychologists realized that the presence of other people is *socially arousing*; that is, people become more aware both of their own behavior and of social norms when they are in groups. Social arousal *facilitates well-learned habits, but it inhibits poorly learned habits*. Thus, when trying to teach employees new habits to improve security, it is counterproductive to put them into large groups. Individualized learning (e.g., by means of computer-based training and videotapes) can overcome inhibitory effects of groups in the early stages of behavioral change.

50.5.2 Locus of Control. Another factor that interferes with implementation of security policies is the *locus of control.* People do not like feeling that they have no control over their environment. For example, in a classic experiment reported in social psychology textbooks, two equivalent teams of people were both subjected to loud and disruptive noise coming through a loudspeaker in their work area. One group had no control whatever over the noise, whereas the other had a large button with which they could stop the noise at once. The group with the stop button did noticeably better at their complex task than the other group—yet in no case did anyone actually press the button. Simply feeling that they *could* exert control, if they wanted to, significantly altered the performance of the experimental subjects.

Similarly, in studies of healing among older patients, three groups were defined: (1) controls, (2) people given a plant in a pot, and (3) people given a plant in a pot plus instructions to water it regularly. The third group did significantly better than the second in their recovery. Once again, the sense of control over the environment appeared to influence outcomes.

In security policy implementation, experience confirms that those organizations with the most participation and involvement by all sectors do best at developing and implementing information protection plans. A common phrase that refers to this phenomenon is *buy-in*, as in: "The different departmental representatives felt that they could genuinely buy into the new policies because they had fully participated in framing them."[15]

50.5.3 Group Polarization. Another branch of research into group psychology deals with group polarization. Groups tend to take more extreme decisions than would individuals in the group acting alone. In group discussions of the need for security, polarization can involve deciding to take more risks—by reducing or ignoring security concerns—than any individual would have judged reasonable. Again, one-on-one discussions of the need for security will generally be more effective in building a consensus that supports cost-effective security provisions than will large meetings.

50.5.4 Groupthink. In the extreme, a group can display groupthink, in which a consensus is reached because of strong desires for social cohesion. When groupthink prevails, evidence contrary to the received view is discounted; opposition is viewed

as disloyal; dissenters are discredited. Especially worrisome for security professionals, those people in the grip of groupthink tend to ignore risks and contingencies. To prevent such aberrations, the leader must remain impartial and encourage open debate. Respected security consultants from the outside could be invited to address the group, bringing their own experiences to bear on the group's requirements. After a consensus—not the imposition of a dominant person's opinions—has been achieved, the group should meet again and focus on playing devil's advocate to try to come up with additional challenges and alternatives.

In summary, security experts should pay attention to group dynamics and be prepared to counter possible dysfunctional responses that interfere with acceptance of IA policies.

50.6 TECHNOLOGICAL GENERATION GAPS. In our society there are growing societal gaps between the social groups that grew up interacting in real-world communities (unwired), groups that grew up with the Internet (wired), and the newest group growing up with the always-on technology of our complex and content-rich wireless social networks.

- **The unwired generation.** In today's always-on world of ubiquitous wireless communications, we sometimes forget about the unwired generation, those born in the early 1960s or before, who grew up actually playing outside with their friends and communicating face to face. For these older employees, online discussion groups, streaming video training, blogs, and email may not be as effective as with other, more technological groups. The unwired generation may see the use of impersonal technology in training as an indicator of management apathy toward a topic that is not worth taking the time for real-world interaction. To this generation, if security training is important, someone should take the time to deliver it face to face. The unwired generation is also the most susceptible to many of the social psychology techniques and pitfalls discussed in this chapter.
- **The wired generation.** Those born from the early 1960s to the late 1970s comprise the wired generation. This transitional generation grew up at the dawn of the Internet from 300 baud dial-up access to ISDN. From MS DOS to Windows 98 and from Cobol to C++, this generation is the bridge from the real world to the cyberworld. This generation is also the generation currently coming to power in both business and government. Although accustomed to meeting face to face for important business, the wired generation lives by email and cell phones, tolerating both the unwired and wired methods of communicating and learning.
- **The always-on generation.** Those born after 1980 comprise the always-on generation. This generation grew up with high-speed Internet, cell phones, video games, portable electronics, and online virtual communities. This is the connected, networked, Facebook, video-game generation. It is this generation that seems to confound traditional security screening and implementation policy.

Although the unwired and wired generations shared psychological roots in the real world, the always-on generation has two homes: the real world and the cyberworld. With instant messaging (IM), texting (Short Message Service, or SMS), and cellular wireless, this generation moves in and out of cyberspace as the previous generation moved between the worlds of work and home. Moreover, like the previous generations, the always-on generation has different psychological and sociological frameworks for their worlds—a serious information assurance concern.

To the unwired and wired generations, personal communication meant face to face. To the always-on generation personal, communication means IM, email, and smart phones. In addition, social interaction often means social networking, blogs, Web pages, message boards, and wikis. Add to this the fact that the always-on generation has a keen ability to quickly transfer information between their tightly integrated worlds, and you have the ingredients for a security officer's nightmare.

Recognizing these differences and reacting accordingly can pay sizable dividends in both security compliance and general management success. As a first step, one can ensure that security teams include members from more than the oldest generation in the enterprise; younger people may be able to act as intermediaries or translators between increasingly disparate cultures.

50.7 SUMMARY OF RECOMMENDATIONS. This chapter has reviewed the major findings of social psychology that can help to improve IA programs. These ideas can prove useful to readers who think about social psychology as they work to implement security policies:

- Recognize that IA policies often conflict with the schema for trusting, polite behavior in situations outside the work arena.
- Train IA personnel to recognize that failure to comply with security policies may be rooted in many other factors than simply bad attitude.
- Listen more than you command.
- Teach security personnel to avoid the classic errors of attribution when trying to understand their colleagues' motivations.
- Openly discuss and counter prejudice before it causes conflicts.
- Take intercultural differences into account when setting and implementing security policies.
- Before attempting to implement policies and procedures, ensure a consistent view of IA among colleagues.
- Whenever possible, security policies should be introduced over a long time, not rushed into place.
- Presenting case studies is likely to have a beneficial effect on participants' readiness to examine security requirements.
- Security awareness programs should include many realistic examples of security requirements and breaches.
- Attempt to inspire a commitment to security rather than merely describing it.
- Emphasize improvements rather than reduction of failure.
- Create a new concern for corporate security by exploring the current structure of beliefs among employees and managers.
- Never portray computer crime using positive images and words.
- Praise any comments that are critical of computer crime or that support the established security policies.
- Employees who dismiss security concerns or flout the regulations should be challenged on their attitudes, not ignored.
- Identify the senior executives most likely to succeed in setting a positive tone for subsequent security training and engage their cooperation to act as role models.

- Examples of frightening consequences used in awareness and training materials should be coupled immediately with descriptions of effective and achievable security measures to forestall such consequences.
- Presenting objections to a proposal and offering counterarguments is more effective than one-sided diatribes.
- Security-awareness programs should include many, frequent, and preferably novel and entertaining reminders of security issues.
- In addition to tapes and books, rely on a charismatic teacher or leader to help generate enthusiasm for better security.
- Encourage specific employees to take on public responsibility for IA within their work groups.
- Rotate security roles periodically.
- Security training should include information on how to tell that someone may be engaging in computer crime.
- Build a corporate culture that rewards responsible behavior, such as reporting security violations.
- Develop clearly written security policies and procedures.
- Security procedures should make it easy to act in accordance with security policy.
- Treat failures to act in accordance with security policies and procedures as very serious matters.
- Enforcing standards of security can increase the likelihood that employees will subsequently act more cooperatively.
- A working environment in which employees are respected is more conducive to good security than one that devalues and abuses them.
- Get to know the staff from whom you need support.
- Encourage social activities in the office.
- Pay special attention to social outliers during instruction programs.
- Monitor compliance more closely at both ends of the popularity range.
- Work with the outliers to resist a group's antisecurity bias.
- Include small gifts in your security-awareness program.
- Start improving security a little at a time, and work up to more intrusive procedures.
- Before discussing security at a meeting, have one-on-one discussions with the participants.
- Remain impartial, and encourage open debate in security meetings.
- Bring in experts from the outside when faced with groupthink.
- Meet again after a consensus has been built, and play devil's advocate.
- Recognize the generational technology gaps in our culture and communicate accordingly; therefore, include people from different generations in your security teams.

None of these suggestions is essential; none of them is appropriate in all situations. However, building on the accumulated experience and wisdom of social psychologists will support the smooth integration of information assurance into any corporate culture. We hope that readers will explore the literature of social and organizational psychology

and will try out new ideas that will enrich the field of information assurance in years to come.

50.8 FURTHER READING

Adler, N. J., and A. Gunderson. *International Dimensions of Organizational Behavior*, 5th ed. South-Western College Publications, 2007.

Bik, O. *The Behavior of Assurance Professionals: A Cross-Cultural Perspective.* Eburon Academic Publishers, 2010.

Greenberg, J. *Managing Behavior in Organizations*, 6th ed. Prentice-Hall, 2012.

Kreitner, R., and A. Kinicki. *Organizational Behavior*, 9th ed. McGraw-Hill, 2009.

Kowert, P. *Groupthink or Deadlock: When Do Leaders Learn from Their Advisors?* State University of New York Press, 2002.

Lesko, W. A. ed. *Readings in Social Psychology: General, Classic, and Contemporary Selections*, 8th ed. Pearson, 2011.

Mills, J. H., K. Dye, and Al. J. Mills. *Understanding Organizational Change*. Routledge, 2008.

Myers, D. G. *Social Psychology*, 11th ed. McGraw-Hill Humanities, 2012.

Myers, D. G. *Exploring Social Psychology*, 7th ed. McGraw-Hill Higher Education, 2014.

Prestwich, A., and M. Conner. *Applied Social Psychology.* Wiley-Blackwell, 2014.

Senior, C., and M. Butler. *Social Cognitive Neuroscience of Organizations*. Wiley-Blackwell, 2008.

Smith, E. R., and D. M. Mackie. *Social Psychology*, 3rd ed. Psychology Press, 2007.

50.9 NOTES

1. This chapter is based on original work by M. E. Kabay as a contributed paper at the Sixteenth National Computer Security Conference organized in 1993 by the National Computer Security Center. That work was updated over the years and became a chapter in the fourth edition of this *Handbook.* It was updated for the fifth edition with contributions from colleagues teaching and studying in the MSIA (now MISA) program at the School of Graduate and Continuing Studies at Norwich University. This version includes minor updates and new suggested readings.
2. Few specific references to the scholarly literature of social psychology research are included in this chapter except for quoted materials. For details of the information presented, consult any college-level introduction to social psychology.
3. J. Baron, *Thinking and Deciding*, 4th ed. Cambridge University Press, 2007
4. Government Accountability Office, "Information Security: Emerging Cybersecurity Issues Threaten Federal Information Systems," United States Government Accountability Office Report GAO-05-231 (May 2005); www.gao.gov/new.items/d05231.pdf
5. M. E. Kabay, "The VA Data Insecurity Saga." (2008), www.mekabay.com/infosecmgmt/vasaga.pdf
6. R. D. Austin and C. A. R. Darby, "The Myth of Secure Computing," *Harvard Business Review* (June 2003); http://cb.hbsp.harvard.edu/cb/web/product_detail.seam;?E=71608&R=R0306J-PDF-ENG

7. For an in-depth discussion of the manager's role in IA, see Chapter 63 in this *Handbook*; for a discussion of the role of the chief information security officer, see Chapter 65.
8. N. Tal-Or and Y. Papirman, "The Fundamental Attribution Error in Attributing Fictional Figures' Characteristics to the Actors." Paper presented at the annual meeting of the International Communication Association, Sheraton New York, New York, NY, April 13, 2008; www.allacademic.com/meta/p13476_index.html
9. G. Hofstede, G. J. Hofstede, and M. Minkov, *Cultures and Organizations: Software of the Mind,* 3rd ed., McGraw-Hill, 2010.
10. S. Salahuddin, "Taliban Defend Yellow Badges for Non-Muslim Afghans," Reuters, May 23, 2001; www.afghanistannewscenter.com/news/2001/may/may23c2001.html
11. See Chapter 45 in this *Handbook* for a discussion of such an example.
12. For more ideas on effective security-awareness programs, see Chapter 49.
13. SIIA Resource eStore: www.siia.net/estore/10browse.asp and choose Category= Anti-piracy
14. See FUTURECents, "Management by Walking Around, www.futurecents.com/mainmbwa.htm; and A. Fisher, "Management by Walking Around: 6 Tips to Make it Work," CNNMoney, August 23, 2012, http://management.fortune.cnn.com/2012/08/23/management-by-walking-around-mbwa/
15. See Chapter 66 in this *Handbook* for a discussion of the importance of widespread participation in security-policy development.

CHAPTER 51

SECURITY STANDARDS FOR PRODUCTS

Paul Brusil and Noel Zakin

51.1 INTRODUCTION **51·1**

51.2 IMPORTANCE OF STANDARDS **51·2**

51.2.1 Value of Standards 51·2
51.2.2 Motivation for Establishing Trust and Managing Risk in Products 51·3
51.2.3 Motivation for Establishing Trust in Systems 51·4

51.3 TYPES, SOURCES, AND EXAMPLES OF STANDARDS **51·4**

51.3.1 Types of Standards 51·5
51.3.2 Sources of Standards 51·5
51.3.3 Examples of Security Standards 51·7

51.4 PRODUCT DEVELOPMENT AND ASSESSMENT APPROACHES **51·9**

51.4.1 Historical, Sunsetted Approaches 51·10
51.4.2 Consensus Security Specification Approaches 51·10
51.4.3 Standard Product Development Approaches 51·11
51.4.4 Informal Product Assessment Approaches 51·14
51.4.5 Security Assessment Standards for Products 51·22

51.5 STANDARD PRODUCT AND SERVICES ASSESSMENT APPROACHES **51·23**

51.5.1 Government Standard Cryptographic Validation Programs 51·23
51.5.2 FedRAMP 51·24
51.5.3 Common Criteria 51·26

51.6 NOTES **51·32**

51.1 INTRODUCTION. Organizations, people, devices, tests, software, and products have defining characteristics. Standards are established for the purpose of providing uniformity of essential characteristics among different instantiations of such entities. Standards thus allow different parties to understand and to compare the defining characteristics of different entities.

There are many types of standards depending on the characteristic being addressed. There are also many different bodies or authorities that develop and issue standards.

A particular goal of standardization in the security arena is to evolve toward an information technology (IT)–driven economy where security products, and secured

products, approach plug-and-play status. They should be comparably trusted, be available for purchase from multiple competing vendors, and be able to be mixed, matched, and integrated to provide requisite secure, trusted IT infrastructures that reduce the risks of greatest concern.

In the U.S. Government sector, product security and software quality assurance has been important for years. However, security became of top concern to government agency chief information security officers as a result of the *Federal Information Security Management Act* (FISMA).[1] FISMA applies to the security of systems.

The security of *systems* such as those impacted by FISMA and other industry-specific standards applicable to secure systems is not considered in this chapter. But such standards have certainly increased the attention to security associated with an individual *product*.

Although several different categories of security-relevant standards exist, this chapter provides insights about the field of security standards by considering only a representative sample of the multitude of standards that pertain to product security. Standards associated with developing product trust are especially important in the electronically interwoven world. As such, the primary focus herein is on standards pertaining to product trust.

The chapter begins by providing a general introduction to standards. Section 51.2 addresses why standards are important. Section 51.3 summarizes what types of security-relevant standards exist, what bodies create standards, and what security characteristics, features, or capabilities are addressed in examples of different standards.

Then attention turns to considering examples of some of the many types of standards that apply specifically to enhancing trust in products. Section 51.4 describes several standards-based product development approaches ranging (for example) from consensus-based standards for security specifications and product development processes to formal capability-based standards and ISO secure software development processes. Section 51.4 also summarizes several different product assessment approaches that rely on informal standards ranging from vendor self-declarations to third-party assessments via *de jure* methods.

In Section 51.5 three standards-based approaches for assessing the security of products are described, for products ranging from cryptographic modules and products, to general classes of products, to cloud service offerings.

51.2 IMPORTANCE OF STANDARDS.

The value of standards is considered in Section 51.2.1. The motivation for using standards to establish trust and to manage risks associated with products is discussed in Section 51.2.2. The difficulties of establishing trust in systems consisting of many individual products are briefly examined in Section 51.2.3

51.2.1 Value of Standards.

Many parties benefit from standards pertaining to security: customers, vendors, software developers, testing houses, product reviewers in the media, consultants, and more.

For example, customers find standards helpful in several ways. Standards help specify their needs for various security functionalities and the degrees of assurance they require in the products they buy. Standards help customers understand what security functionality and assurances a product builder claims to provide. Standards help consumers select commercial off-the-shelf products that they can trust will conform to their security and assurance requirements and that, as needed, will interoperate with comparable products. Procurement standards help customers buy appropriately

trusted products that meet customer requirements. Customers under the mandates of the security-relevant regulations imposed by federal regulations such as the *Health Insurance Portability and Accountability Act* (HIPAA) and the *Sarbanes-Oxley Act* (SOX) often look to establishing due diligence by leveraging products that have established trust in their security and assurance functionality in a standard way. Standard security certifications of customer staff and hired consultants help customers employ appropriately knowledgeable people to protect customers' facilities and information or to help deal with certain vendors or regulators.

Vendors and service providers find standards helpful in several ways. Use of standards provides evidence that vendors have migrated their product development to a paradigm wherein security is built in from the start. Use of standards provides evidence that security is not some afterthought that is patched or bolted on. Use of implementation-related standards and standard personnel certifications shows that security is the foundation on which a vendor is building a product or a service provider is providing a service. Use of standards helps to open global marketplaces to vendors and service providers. By using standard, third-party verification of security capabilities, vendors are making their products (and service providers are making their services) either stand out from, or be comparable to, their competitors.

51.2.2 Motivation for Establishing Trust and Managing Risk in Products.

Trust in the electronic processing, storage, and interactions among customers, businesses, business partners, suppliers, service organizations, and governments is key to electronic economy, electronic healthcare, and electronic government models. The need for trust will only increase as new IT paradigms and technologies proliferate, mutate business and IT support models, and introduce new risks and vulnerabilities. For electronic service models to succeed, all e-business, e-health, and e-government players need confidence in the IT products used by interacting players as well as products used in the intervening IT infrastructure.

Another key notion besides trust is the notion of risk management. When electronic relationships are established between parties, there are quantifiable risks associated with such relationships.

Risks are quantifiable in many ways. For example, they can be quantified in terms of the types of possible adverse events. They can be quantified in terms of the likelihood of different types of adverse events and by the value of what is to be protected by IT security solutions during an adverse event. They can be quantified in terms of the consequences of adverse events, such as the liability that may be exposed via compromises, or the entities that may be hurt by compromises.

Risks then can be mitigated in a number of possible ways. For example, risks can be mitigated by using products that reduce the occurrence or impacts of the adverse events of most concern. When assets of increasing value need to be protected, risks can be reduced by using products that have increased assurance. Risks also can be mitigated by using products that decrease the specific, deleterious liabilities and undesired consequences of greatest concern. Being able to specify the risks of concern and to identify security solutions that mitigate those risks is a powerful strategy used within certain standards-based approaches such as the *Common Criteria* (CC) paradigm.

Linking notions of trust together with notions of risk reduction makes for a powerful strategy. Establishing trust among interacting electronic parties that all parties are using products with appropriate security quality is akin to the role of face-to-face handshakes in older business models. In the e-business marketplace, trust, like the old handshake,

is key to increased revenues via increased business transaction volumes. Furthermore, by mitigating risks, business losses and costs can be reduced. When trust enhancement is coupled with risk management, the resulting increased revenues, combined with decreased losses and lower costs, make for significant profit multiplication.

The CC paradigm is an example of a standards-based product assessment strategy that provides for both establishing appropriate levels of trust in products and for specifying, managing, and mitigating risks of most concern.

Various standard-based ways have been developed and used over the years to build confidence about the quality of security implementations. The focus of such efforts includes establishing trust via one or both of two perspectives: (1) that a product performs its claimed security functionality completely and correctly and (2) that the product builder's processes (from design, to development, to delivery, to maintenance) are sound. Typically, trust is established either by testing in a standard way just the implementation or by evaluating in a standard way both the product and its implementer.

Not all approaches to establishing trust via product assessment necessarily address trust via formal standards–based assessment of implemented security functionality as well as the assurance of the soundness of the builder's abilities and processes. That is, many of the testing and evaluation approaches are based on informal, de facto or de jure standards[2]—not on formal, nationally recognized standards.

51.2.3 Motivation for Establishing Trust in Systems. Recent U.S. Government administrations have recognized the importance of establishing trust in IT. They have issued directives and guidance to elevate awareness of the central, critical nature of IT and to help to preserve trust in national IT infrastructures.

National IT infrastructures are typically large IT systems often composed of many individual products. Standards-based methods for establishing trust in products do not necessarily work for establishing trust in systems. Standards-based methods of establishing trust in IT systems exist. They include, for example, standards and testing activities associated with the original U.S. *Department of Defense* (DoD) *Information Technology Security Certification and Accreditation Process* (DITSCAP)[3] and the *Department of Defense Information Assurance Certification and Accreditation Process* (DIACAP)[4] which replaced it, the U.S Government's FISMA, and the financial community's *Payment Card Industry Data Security Standard* (PCI DSS).[5]

Such systems-focused, trust-development methods tend to be employed only by large enterprises. Furthermore, the problem of establishing a quantitative measure of trust in a very large, heterogeneous system like a national IT infrastructure is evolving. At a minimum, it requires coordination and cooperation among all who contribute to or use the infrastructure. Security requirements and security testing approaches applicable to IT systems, per se, are not considered herein.

Although still argued among security professionals, many believe a step in the right direction is to build IT systems with products and components that are individually assessed to be trustworthy, with some specified degree of confidence. What is most difficult, however, in a product assessment is to test the product's security aspects in an environment that perfectly mimics the environment in which the product is to be used.

51.3 TYPES, SOURCES, AND EXAMPLES OF STANDARDS. In today's heterogeneous, multidisciplinary, multitechnology, multiparty, interconnected IT environments, different types of security standards are essential for a variety of reasons. Several formal, informal, and ad hoc bodies develop and issue standards. The significance and acceptance of any given standard tends to depend on the notoriety of

the body that creates the standard. A high-level taxonomy for what types of security standards exist appears in Section 51.3.1. The various types of bodies involved with security standardization efforts are listed in Section 51.3.2. Examples of the standards developed and issued by various bodies are presented in Section 51.3.3

51.3.1 Types of Standards. In the security world, there are a many different types of standards. The types of standards pertinent to security include, for example, the following:

- Capability standards
- Personnel certifications
- Risk assessment criteria
- Requirements specifications
- Functional specifications
- Assurance specifications
- Performance criteria
- Product development standards
- Testing, evaluation, and assessment standards/criteria[6]
- Product review criteria
- Interoperability standards
- Procurement standards
- Ancillary standards

51.3.2 Sources of Standards. Some of the bodies that issue formal nation-recognized standards, organizational standards, or proprietary standards, or that use de jure approaches pertinent to security, include the following:

- Recognized national bodies
- International organizations
- Governments
- Military
- Consortia
- Like-minded communities
- Vendors
- Consulting houses
- Trade press
- Commercial testing houses
- Private testing contractors

Formal standards relevant to information assurance (IA) are created, published, and maintained by recognized standards bodies. There are various recognized standards bodies including technology-specific working groups associated with professional organizations like the IEEE (Institute of Electrical and Electronics Engineers) and the IETF (Internet Engineering Task Force). Standards are also created by working groups

associated with recognized, country-specific, national standards bodies, such as the ASC (American Standards Committee, formerly ANSI) in the United States, the BSI (British Standards Institute) in the United Kingdom, and the *Bundesamt für Sicherheit in der Informationstechnik* (also BSI) in Germany. Such national standards bodies create either country-specific standards or they collaborate and harmonize with other peer national standards bodies to create international, globally applicable standards, such as those associated with International Organization for Standardization (ISO).[7]

Recognized organizations within national governments also create standards. For example, the National Institute of Standards and Technology (NIST)—a U.S. federal agency under the U.S. Department of Commerce—creates and issues standards called *Federal Information Processing Standards* (FIPS). It also issues standards called SPs (*Special Publications*) that delineate security, privacy, and assurance best practices guidance. FIPS and SPs often apply not only to the U.S. Government but are widely used as references in other nongovernmental organizations and may even apply formally to some. For example, a 2005 amendment to the *Federal Acquisition Regulation* to implement the IT security provisions of FISMA requires all federal procurements to adhere to pertinent NIST standards and publications.

National governments also form bodies to collaborate and to harmonize international standards for use by governments within several nations. The CC Project is an example of such a cross-government standardization body. Standards developed by this particular cross-government standards body are also fed to recognized international standards bodies such as the ISO to create even broader-reaching standards.

In the military area, the U.S. DoD publishes military standard directives such as *Mil Standard 8500*,[8,9] that must be followed when procuring security-related products, as well as *Department of Defense Directive 8570*[10] that pertains to training, certification, and management of employees performing IA functions.

Technology-specific and industry-specific consortia have large, broad-based, international membership; examples include

- Cloud Security Alliance (CSA)[11]
- Security Standards Council (SSC)[12] of the Payment Card Industry (PCI) consortium
- Software Assurance Forum for Excellence in Code (SAFECode)[13]
- Object Management Group (OMG)[14]
- Web Application Security Consortium (WASC)[15]
- TM (formerly TeleManagement) Forum[16]

Often they are considered to be creators of informal, security-relevant standards for use within specific industries or technology sectors. Small consortia or individual entities also claim to create standards, but such efforts are generally not accepted by as broad constituencies as are the legitimate, recognized, formal, and informal standards.

Other like-minded communities like the *SANS* (formerly *System Administration and Network Security*) *Institute*[17] promulgate de facto best-practices security standards, such as the *20 Critical Controls*,[18] claimed to reflect community consensus within a certain community of security practitioners.

Vendors such as Microsoft develop standard-like documents such as the *PCI DSS Compliance Planning Guide*[19] or the *Secure Development Lifecycle*[20] framework

and strategy for secure coding that they presumably hope will be adopted by large communities.

Consulting houses such as Gartner, Forrester, and Frost & Sullivan develop private product review criteria pertaining to security and security features. They use such criteria to evaluate security-relevant products. Gartner and others also produce product reviews based on user interviews.

The trade press includes media outlets such as *Computerworld, CSO, IT Business Edge, Network World, SC Magazine*, and *Security Management* that create yet other private product review criteria that they use to examine security-relevant products.

Commercial testing houses such as *ICSA Labs*, *NSS Labs*, and *West Coast Labs* create and use private, independent, product-testing criteria and methodologies to evaluate product security and performance. Some of the commercial testing houses, such as AV-Test, develop private, independent, product-testing criteria and methods to certify tested products.

Private contractors conduct product security testing via a variety of informally standard techniques that rely on private criteria and test procedures. Such techniques include vulnerability scanning, vulnerability assessment, penetration testing, security auditing, ISO 27001 auditing, security reviews, ethical hacking, and more. Academic organizations are also beginning to establish security product-evaluation laboratories that rely on various published vulnerability and assurance test procedures and criteria.

51.3.3 Examples of Security Standards. Examples of the types of security-relevant standards issued by examples of the different standards issuing bodies are provided in this section below. Detailed summaries of a number of these standards are provided later in this chapter.

Capability standards such as the Capability Maturity Model (CMM) and ISO 9000 exist to measure the competency of organizations that build products—in particular, security-related products.

General personnel certifications such as the Certified Information Systems Security Professional (CISSP) from the International Information Systems Security Certification Consortium ((ISC)^2) measure, among other things, the competency of staff specifying, building, procuring, installing, or maintaining security-related products. Other certifications like the IEEE Computer Society's Certified Software Development Associate (or Professional) or the (ISC)^2 Certified Secure Software Lifecycle Professional (CSSLP) or the Electronic Commerce Council Certified Secure Programmer are examples of certifications more directly useful in certain product software development areas involving security. The Cloud Security Alliance's Certificate of Cloud Security Knowledge assesses an individual's competency in major matters related to cloud security.

Standards exist to provide consistent ways to stipulate security needs and requirements in both security products and secured products. Some standards specify the security functionality appearing within products. Standards like IETF's IPSec not only specify security functionality, but they also foster interoperability of separately built security implementations. Other standards specify security-related software interfaces, naming conventions, and data structures such as for Common Object Request Broker Architecture (CORBA) middleware products.

Standards like FIPS 140 can specify security requirements, evaluation methods, and validation concepts for a single specific type of security-relevant component or product. The CC series of standards can specify the security functionality and security assurance characteristics to be incorporated into any class of products.

De jure standards like SANS 20 Critical Controls, which compete with certain national and international standards, specify a single, small, essential core of security functionality. This de jure, consensus-based standard may be growing in popularity.

Product development standards like ISO/IEC 27034-1 specify requirements for secure software development processes while the Building Security in Maturity Model (BSIMM) de facto industry standard provides an ability to compare secure coding environments and to identify best development practices.

Vendor standards like Microsoft's Security Development Lifecycle also specify software development and assurance processes to increase reliability of software security.

Standards also exist to govern the testing of products. It is essential that vendors and buyers of security and security-enhanced products (a) verify that a product under test implements functionality that addresses the stipulated security functional requirements desired, and (b) demonstrate how well the product and the vendor's product development processes meet stipulated requirements. There are national and international standards such as FIPS 140 and the CC that stipulate such conformance testing and that stipulate testing methods to accomplish such testing. The FedRAMP standard approach applies to validation of the security capabilities associated with a specific class of service providers.

Conformance of different security-relevant products to the same functional standard, however, does not necessarily ensure that these products will interoperate. Interoperability of separately manufactured security solutions is possible if implementation standards can specify precise, interoperable profiles of selected security capabilities from within a security functional standard. Interoperability testing can assure secure interoperation between comparable products built by autonomous vendors and used between autonomous parties. Conformance implies only that interoperability is possible; but interoperability needs to be verified by pair-wise testing. Standards exist to specify how to examine conformance of implementations to functional standards and how to assure the interoperability of implementations that must meet the same functional standard.

Interoperating, conformant, secured products and security products cannot necessarily be trusted to provide or to support sound security, or to mitigate the risks of greatest concern. Key to developing trust is to build confidence that products mitigate the risks of concern, that products are properly built and behave according to specification, and that products do no more or no less than advertised. Standards like the CC exist for establishing and testing the degree to which risks and vulnerabilities are mitigated to a specified level of confidence. These standards specify implementation assurance requirements and associated tests and testing methods to verify the quality of implemented security within the products under test.

Regarding private, de jure standards, different trade press media and others develop private product review criteria and methodologies. These private standards are used to create reviews of specific security or security-enhanced products, or to develop buyer's guides that compare the individual product reviews of several products within a class of products. Certain metrics created by such criteria can provide measures of security performance.

Some media outlets operate awards programs that use question-driven surveys for identifying readers' choice picks of best products within certain classes of products. Popularity-based product leadership awards also exist.

In the area of procurement standards, standards like DoD's *8500.1* and *8500.2* stipulate standard policy and policy implementation directives, respectively, for what types of products DoD must procure. These standards also specify what must be done

by DoD entities to meet a need or what constraints DoD organizations must follow when procuring IT products.

In the area of ancillary standards, standards such as companion standards within the suite of CC standards specify how to become a product tester and how to accredit organizations to conduct conformance, trust, or interoperability testing in a standard way.

Other ancillary standards—not considered in this chapter—like the DITSCAP[21] (replaced by DIACAP in November 2007) for U.S. DoD, the National Information Assurance Certification and Accreditation Process (NIACAP),[22] and the NIST Certification & Accreditation (C&A)[23] methodologies for U.S. federal civilian departments and agencies specify how to certify and how to accredit systems composed of secure products and secure components. Standards like ISO 17799[24] and NIST SP 800-53[25] can specify guidelines or best practices for users assembling and using secure components.

In aggregate, standards promote consistent security, end-to-end as well as across different public and private domains and computing environments.

51.4 PRODUCT DEVELOPMENT AND ASSESSMENT APPROACHES.

This chapter focuses on standards for enhancing trust in security-related products. There are two classes of products of interest: "security" products and "security-enabled" products. Security products directly provide security services or prevent penetrations. Security products include, for example, intrusion detection products and firewalls. Security-enabled products are secured products that do not exist solely to provide security services; instead they provide other services that are protected. Examples include operating systems, database management systems, and virtual private networking gear that incorporate security functionality such as identification and authentication or IPSec to protect either the product or the services provided by the product.

The remainder of this chapter summarizes various approaches that support increasing trust in such products. The historical approaches for establishing trust in certain classes of products are summarized in Section 51.4.1. A sampling of community-based consensus standards that specify security controls are described in Section 51.4.2. The chapter then summarizes numerous other current approaches for developing product trust, including standard product development approaches used during the design and building of products (Section 51.4.3), informal product assessment criteria and approaches (Section 51.4.4), and formal, internationally recognized, standard product assessment criteria and approaches (Sections 51.4.5 and 51.5).

The U.S. Government had hoped that all public sector and private-sector consumers, vendors, and testers of security-related products would rally around the formal, internationally recognized standards that pertain to specification, development, and testing and evaluation of such products. With formal standards being produced by national and international standards bodies, it was anticipated that consumers could understand with greater certainty the security and assurance features offered by a product or a service provider. It was hoped that the confidence obtained by using universal, internationally accepted standards would contrast sharply with nonstandard approaches that, because of their sheer numbers of competing approaches, would confront and befuddle consumers.

In today's actuality, the government and military sectors frequently rely on the formal standards-based approaches. But, in contrast—despite the expectations associated with such formal standards-based approaches—the entire marketplace has not necessarily coalesced around the formal standards approach.

Many informal private product assessment alternatives that do not rely on formal standards came into being and are being used with greater or less frequency than the alternatives based on formal standards. There is no shortage of alternative approaches and alternative frameworks on which such approaches are based. Correspondingly, there is no shortage in approaches stipulated in user procurement directives. The marketplace is fractured among many alternatives. A sampling of these alternatives is examined below.

51.4.1 Historical, Sunsetted Approaches. To introduce consistency in describing the security features and levels of trust of a limited set of security-enhanced products, and to facilitate comprehensive testing and evaluation of such products, the U.S. DoD developed the *Trusted Computer System Evaluation Criteria* (TCSEC).[26] The TCSEC—often called the *Orange Book*—defined a small set of classes (C1 to A1) of increasing security functionality and increasing assurance applying to operating systems. The TCSEC was extended to networking devices[27] and database management systems.[28] Government in-house evaluations were offered first, followed by comparable government-sponsored commercial evaluation services.

Partly because of large testing delays and costs, other countries developed other criteria that were more flexible and adaptable to accommodate rapidly evolving IT. The *Information Technology Security Evaluation Criteria*[29] (ITSEC) arose from the combined inputs of earlier German criteria, French criteria, and U.K. confidence levels. The *Canadian Trusted Computer Product Evaluation Criteria* (CTCPEC) were then developed as a combination of the TCSEC and ITSEC approaches. The U.S. Federal Criteria development then attempted to combine the CTCPEC and ITSEC with the TCSEC.

With growth of the international market for trusted IT products, all these historical, competing, national criteria had the potential to fracture the marketplace. They were sunsetted after efforts were completed to harmonize the various criteria into common, standards-based, internationally accepted criteria.

The result of the harmonization effort was a single, wide-ranging CC program (Section 51.5.3). These standard criteria provide a fully flexible, highly tailorable approach to the standardization of security functionality, evaluation assurance requirements, specification, and testing of implementations of security-related products.

51.4.2 Consensus Security Specification Approaches

51.4.2.1 SANS Institute. The SANS Institute has issued a de jure standard[30] that specifies what is believed to be the core set of the most critical, baseline security controls deemed essential for virtually any organization. These security controls are colloquially called the *20 Critical Controls*. They provide a feasible, implementable, best-practice, security approach based on providing what is purported to be the greatest risk-reduction, the best protection against actual, in-the-field threats, and the best techniques to stop the most dangerous, common, regularly occurring attacks. The SANS approach provides for continuous, automated monitoring of key portions of an organization's IT infrastructure to ensure security effectiveness.

SANS has published a guideline of the steps for how to implement these critical controls.[31] The guideline provides a tabulation of many real attacks and provides a mapping that identifies which of the 20 Critical Controls provide the best defenses against each of these attacks.

The SANS Critical Controls map to a certain subset of the security controls specified in NIST Special Publication 800-53, Revision 3 (see Appendix A of the SANS Critical Controls guidelines). In some sense, the SANS critical controls may be considered as a specific profile of the most essential security controls in NIST SP 800-53. Some security professionals feel the SANS 20 Critical Controls will provide the necessary (but not necessarily sufficient) due diligence to become the de facto yardstick for measuring the standard of care for security.

51.4.2.2 Cloud Security Alliance (CSA). The CSA has developed a guideline explaining cloud computing and identifying the best risk-based practices for transitioning to secure cloud operations provided by different cloud service models in cloud computing environments. Partitioning of security responsibilities between cloud customers and cloud service providers depends on the cloud service being offered, for example, cloud Software as a Service (SaaS), cloud Infrastructure as a Service (IaaS), cloud Platform as a Service (PaaS). The stipulated security controls and technologies apply to (a) protecting data migrating to clouds, (b) protecting data in transit to the cloud and in transit between different cloud providers/cloud environments, (c) protecting data within clouds, (d) privacy preservation, (e) application security, (f) data loss prevention, and more. The version 3 edition of this guideline is available at the time of this writing (July 2013).[32]

51.4.3 Standard Product Development Approaches. A key that can lead to a productive product assessment is for product builders to rely on high-quality, highly capable product developers and to use product development standards, such as software design and testing standards, during product design and creation. Such product developer capability standards that can apply to builders of security products include the Capability Maturity Model (Section 51.4.3.1) and the ISO 9000 series of standards (Section 51.4.3.2). Standards and approaches used during product development to enhance the quality and inherent security offered by products are examined in Section 51.4.3.3.

51.4.3.1 Capability Maturity Model. The Systems Security Engineering Capability Maturity Model (SSE-CMM) has been standardized as ISO/IEC 21827. This standard provides a way to assess the soundness of a security product builder's engineering practices, as well as a way for continuously improving such practices up a hierarchy of increasingly mature software process levels, during the many stages of product development, such as during:

- Product requirements capture and analysis
- Product concept definition, including accurate translation of security requirements into product requirements
- Product architecting
- Product design
- Product implementation

A security product developer can demonstrate competence in building products by means of recognized, so-called capability maturity assessments of the developer's software and security engineering processes. Security-enhanced products built by organizations with demonstrated expertise and maturity can merit greater trust than products

built by organizations that do not demonstrate mature, competent, software design, and security engineering capabilities.

The SSE-CMM establishes a framework of generally accepted security engineering principles and a standard way of measuring (and improving) the effectiveness of an organization's security engineering practices. The SSE-CMM describes the essential characteristics of, and provides tools for, assessing an organization's security engineering process that must exist to ensure good security engineering. These characteristics are graded by a set of security metrics that assess specific attributes about a vendor's processes, and the security effectiveness of the results of vendor's processes.

When the level of the SSE-CMM security metrics associated with a specific builder shows the builder to have mature security engineering capabilities and effective security engineering practices, then confidence is increased that the builder can build sound security products.

Trust in, and assurance about, a product can be inferred, to some degree, for measurably competent vendors that use sound security engineering processes as assessed by the SSE-CMM. The quantitative comparability of assurance developed via the SSE-CMM approach to the assurance developed via other approaches such as evaluation of assurance requirements stipulated from the CC paradigm is not well understood. For now, it appears possible to assess the assurance of a vendor's capability to build quality products by both the SSE-CMM and CC approaches; but, perhaps both approaches should be integrated to form a more comprehensive assurance assessment model.

51.4.3.2 ISO 9000 Standards. The ISO 9000 standard[33] is used as a guide to conduct a broad, high-level, horizontal assessment of the quality of systems and of the competence of an organization (typically a manufacturing or service organization) across all its facets. Although not specific to organizations that build security products, it does provide some amount of basic information about the potential for quality and repeatability in an organization's ability to meet its mission. In fact, derivative standards such as in the CC, in part inspired by ISO 9000, are used to accredit the quality associated with security testing laboratories.

The ISO 9001 standard has applicability to software development. Although not specific to organizations that build security products, it nevertheless specifies in general an acceptable, minimum quality level for software processes.

51.4.3.3 Product Development Standards and Approaches. There is great need to focus on secure software development as new products are brought to market. The current situation for secure software development is not stellar. Annually, about 5,000 new software vulnerabilities—many exploitable by real-world threats and attacks—are discovered per year. Unfortunately, software sabotage also incorporates intentional, sinister flaws. Studies indicate that up to 99 percent of coded applications contain vulnerabilities. Industrial control, so-called supervisory control and data acquisition (SCADA) software, is equally vulnerable. Indeed, while slightly better than these statics for general applications, even software within security appliances is not immune to software development flaws and is not—as many customers might believe—inherently secure.[34]

To overcome threats to software, security software scientists are trying to move secure software development practices from an art to a science based on standards. Common mitigation strategies include establishing appropriate and rather specific policies and procedures (that should include coding examples) to govern secure software

development. Specific tactics include, for example, reliance on commonly available, standard, checks-and-balances–oriented technologies such as

- Code signing by developers,
- Code check-in audit monitors and audit trails,
- Multiparty mutually monitoring code reviews,
- Automated code and vulnerability scanning with acceptable false-positive/negative detection rates,
- Live attack scanning, and
- Use of antitampering technology for code integrity checking.

Guidelines for secure software development exist.[35,36] Commercial development testing tools such as static and dynamic code analyzers can also be used to find defects and vulnerabilities in code being developed.

The key to secure software development efforts[37] is to establish standard corporate culture and appropriate policies that:

- Put software security developers at the table when software project requirements are established,
- Incorporate security into software by design,
- Incorporate only vulnerability-scanned, third-party software components[38] and use dynamically linked libraries so that the newest, less vulnerable, updated software components automatically replace older, outdated, more vulnerable software,
- Foster collaborative mutually cross-checking software development,
- Incorporate better testing early in, and throughout, the software development life cycle so as to manage security and quality defects as the code is being written, and
- Establish formal processes for upstream feedback aimed at eliminating root causes of software vulnerabilities, defects, and flaws.

Such actions increase code quality before code is released for manufacturing. Higher quality code tends to have fewer security vulnerabilities, smaller attack surfaces, and fewer holes that need to be plugged after software reaches production status. If internal testing resources are limited, secure software testing services can be outsourced to external firms such as Coverity that provides a Development Testing Platform or the NCC Group that provides direct testing services.

Going even further than just focusing on software development science, many are looking toward more formal product development standards that provide repeatable, proactive, product development and testing processes. Certain vendors such as Microsoft (in 2002), via its *Trustworthy Computing Initiative*, and Adobe (in 2009) initiated leading-edge, standard, internal initiatives to incorporate security into their software designs prior to mass production and sales of their software. Microsoft's *Security Development Lifecycle*[39] (SDL) specifies processes for developing secure software.

More recently, the 2011 international standard *ISO/IEC 27034-1* specifies requirements for, and describes processes and practices for, secure software development. Several large software companies are starting to use this standard to start comprehensive programs for necessarily incorporating security into their products right from

the start of the software design stage. This standard is valuable to both buyers and builders of security-related products. Software builders display due diligence in using the standard to create quality, secure code. Buyers understand that compliant software builders are committed to secure development processes.

The *Software Assurance Forum for Excellence in Code* (SAFECode) consortium's training material allows product developers to stay current with the best practices, principles, and methods in secure software engineering and development. Other organizations such as the *Association for Software Testing* are also trying to build a developer community knowledgeable in software development and testing sciences.

The *Building Security in Maturity Model* (BSIMM) program provides an ability for enterprises—as a combined community—to survey its members to discern a common core of secure software development best practices. The BSIMM has created a de facto secure software development standard by comparing the real-life secure coding environments and practices of its membership to identify best secure software development practices. The model is free under a Creative Commons license.

51.4.3.4 Problems in Development Practices. A disappointing observation related to secure development practices is that a 2013 survey[40] indicated that only about a third of product developers are focused on building security into their software; nearly two-thirds are not necessarily focused on security and do not consider using application attack mitigation standard technologies like Data Execution Prevention (DEP) and Address Space Layout Randomization (ASLR). The situation has the potential to worsen as

- New so-called agile software development and testing methodologies emerge,
- Demands for more complex, cloud-based or mobile software rise—often with greater vulnerability densities, and
- Fielded software is upgraded—often by a patchwork of code created by different software developers with differing skills, with differing knowledge of the software being upgraded, and at differing points in times.

Consumers seem to want products with more features; and vendors seem to have little incentives or liability exposures to invest much time or resources into building solid products.

51.4.4 Informal Product Assessment Approaches. A variety of product assessment approaches are not dependent on formal, standard, internationally recognized, security testing and evaluation approaches. Indeed, most have not reached the status of even a de facto standard, but they are considered, nonetheless, for completeness purposes. Such approaches include:

- Vendor self-declarations (Section 51.4.4.1)
- Consumer assessments (Section 51.4.4.2)
- Consortium-based assessments (Section 51.4.4.3)
- Implicit assessments via open-source code development (Section 51.4.4.4)
- Implicit assessments via hacking (Section 51.4.4.5)
- Assessments by magazines and trade press publications (Section 51.4.4.6)
- Third-party commercial assessments using de jure methods (Section 51.4.4.7)

There are shortcomings to most of these approaches. Some are not considered to offer much value. One of the basic shortcomings is that lack of reliance on formal standards makes for difficulty in comparing the product assessment results:

- Of different products,
- For products tested and evaluated via different approaches, and
- For products tested and evaluated by different testing facilities that all purportedly use the same product assessment approach.

51.4.4.1 Vendor Self-Declarations. An approach to establishing trust is based on the notion of vendor self-declarations. A vendor can unilaterally claim that a specific product meets the security needs of a class of customers and that an appropriate amount of customer-desired confidence can be placed in the product's implemented security features. In part, the confidence associated with this approach is implicitly tied to the past reputation of a vendor or to the customer's past experience in dealing with the vendor. If the vendor's reputation or customer's experiences are good, there is some sense that the vendor may have again done an adequate job of implementing security. This approach, however, lacks measurable ways of quantifying the degree of trust that can be associated with a product. It also lacks measurable ways of comparing the relative degrees of trust that can be associated with different products.

An example of such a self-declaration was recently created by the consulting division of a large computer company.[41] In this example, the vendor stipulated what were identified as the most important criteria in selecting a Managed Security Services Provider (MSSP) and revealed why the vendor's services are the best in the marketplace.

This self-declaration approach may have some merit in establishing a sense of continuing confidence in products that have been updated since an earlier version of the product that has undergone rigorous, standards-based security testing and evaluation. If a vendor is known to have good security engineering capabilities—such as can be assessed, in part, by standard Capability Maturity Model approaches (see Section 51.4.2.1)—and if the vendor can provide reasonable evidence as to the nature of the upgrade or revision since the product version that underwent rigorous assessment, then there can be some qualitative (albeit, quantitatively unknown) degree of confidence about the upgraded or revised product. Under these conditions, customers who have innate trust in the vendor can believe that the quality of the changed product is similar to the quality of the version of the product that was formally assessed.

51.4.4.2 Consumer Assessments. It has been reported that less than 20 percent of software consumers are assessing the security attendant with the application software they buy. Product consumers can, however, develop the requisite substantial technical expertise in-house to test and to evaluate specific security-enhanced products directly.

Vendor standard tools are available to help consumers in specific security technology areas. For example, a GFI Software division has developed an *Antivirus Evaluation Guide* that stipulates a repeatable, readily deployable, testing criteria for customers to evaluate vendor antivirus products.[42] For another example, Ixia has developed an attack simulator that a user of security products can run in a safe environment to emulate a range of attack or configuration scenarios so as to evaluate how well the security products in question detect and block certain attacks.[43] For example, attacks by zombie botnets as well as the full life cycle of DDoS attack scenarios can be simulated.

Alternatively, because of lack of in-house expertise or available staff, consumers may hire a consultant to perform limited security assessments via techniques such as penetration testing or source code analysis or binary analysis when source code is not available. Consumers with more resources may contract a private third-party evaluator, such as a commercial security testing provider (such as Veracode, see Section 51.4.4.7.3) or one of the big consulting houses or systems integrators, to do third-party security testing and evaluation. Many of the third-party approaches are private and tend to be unique.

Some financial institutions have used the in-house assessment approach. Financial institutions as a whole are very careful to make sure that products they use are trustworthy. The security, integrity, and soundness of all products and systems supporting financial institutions must be consistent and verifiable. These institutions fear that any breach of IT security anywhere within their systems will result in a loss of confidence in the entire institution, not just in the specific, subverted IT product.

Many financial institutions developed their own internal standard security specifications and evaluation processes as well as an evaluation methodology to quantify, to compare, to approve, and to certify general security aspects of competing products. One of the consequences of this approach to the customer community at large is that it requires substantial, costly duplication of testing infrastructure across the financial industry as well as the costly duplication of testing support efforts by vendors for those products that are candidates for purchase by multiple customers. With each financial entity funding the establishment of its own testing program, the aggregate testing expenses are raised across all entities within the financial industry.

Furthermore, as the volume of financial devices, such as credit card platforms, operating systems, and thousands of applications, continues to increase dramatically, in-house resources are finding it difficult to keep up. Product customers find this kind of do-it-yourself, in-house testing approach to be a tremendous undertaking in terms of development, implementation, legitimacy demonstration, maintenance, and rejustification. They have found it to be expensive, time consuming, resource intensive, hard to maintain, always open to interpretation and to debate, and always in need of justification to regulators and principals in new markets.

51.4.4.3 Consortium-Based Assessment Approaches. Many consortium-directed approaches exist, or have existed, to demonstrate product interoperability, or conformance of a product to stated security features or to specific security technology standards.

In the Internet world, the notion of implementation bake-offs among trial (preproduct) implementations of emerging IETF standards has been a mainstay in the community for quite some time. For example, the IETF initiated a series of *IPSec VPN Interoperability Workshops* upon culmination of the *Internet Key Exchange Version 2* (IKEv2) standard. Vendors of IKEv2-based preproducts gathered in a common testing facility to test the functionality and interoperability of their preproducts against those of their competitors. Initial test scenarios focused on basic functionality and secure tunnel maintenance.

Other consortia use either their own or standard testing approaches. Examples of assessment efforts of an example set of consortia are summarized below. The consortia herein include the *VPN Consortium* (Section 51.4.4.3.1), the *Cloud Security Alliance* (Section 51.4.4.3.2), and the *Smart Card Security Users Group* (Section 51.4.4.3.3).

51.4.4.3.1 Virtual Private Network Consortium. The *Virtual Private Network (VPN) Consortium* developed an approach for demonstrating conformance of a product

to a specific security standard.[44] The VPN Consortium conducts testing of the IPSec and Secure Sockets Layer (SSL) implementations built by its consortium members.

In the early 2000s, the consortium provided three specific profiles of conformance tests of VPN products implementing the IETF's IPSec standard. For each type of test profile, predefined tasks had to be performed successfully against two different reference test gateways. Due to the nonexhaustive set of tests, passing a VPN Consortium conformance test provided only indications that tested products conform, in limited part, to various standard parts of the IPSec standard. Such tests also provided indications that interoperability may be possible with other products that pass the same tests under the same environmental situations.

The consortium has since focused on interoperability testing. It conducts several classes of IPSec interoperability tests: Basic Interoperability, AES Interoperability, IKEv2 Basic Interoperability, IPv6 Interoperability and Authentication with Certificates Interoperability. The types of capabilities being tested, and the profiles by which tested systems are set up for each class of interoperability testing, are specified. The tests help assure VPN users that IPSec systems configured according to the specified profile are generally interoperable with other IPSec systems also configured according to the same profile.

The consortium also provides interoperability testing for profiles of a variety of SSL VPN features in SSL implementations.

The consortium issues logos for member company products that interoperate with at least 75 percent of the other products in an interoperability test. A list of member companies that offer products that pass any of the specific interoperability tests is maintained.[45]

51.4.4.3.2 Cloud Security Alliance. The CSA has established a publically available Security, Trust & Assurance Registry.[46] Any cloud provider can post self-assessments that document the cloud provider's security practices and their compliance to CSA best practices. Cloud providers' self-assessments are documented either (a) according to a CSA standard Cloud Assessment Initiative Questionnaire, or (c) via a CSA Cloud Controls Matrix that details a cloud provider's security concepts and principles in accordance with the CSA's security guidelines (as described in Section 51.4.2.2).

51.4.4.3.3 Smart Card Security Users Group. Although proprietary, in-house approaches were originally used by individual financial institutions to assess thousands of financial IT products, components, and systems, the financial community banded together as a whole in the *Smart Card Security Users Group.* Through this group, financial institutions can use a single standard—the CC—for product assessments and avoid duplication of their individual product assessment efforts. Benefits of such an alliance included:

- Financial institutions can replace their internal, custom product assessment approaches with a common, universally accepted approach.
- They can pool their resources to address common security testing and evaluation needs by using standards-based, CC security specification and testing schemes recognized across all major financial players.
- They can develop profiles of security requirements for the various common elements of smart cards (e.g., chips, operating systems, applications, crypto engines).

- They can develop common test suites to unify the current hodgepodge of fragmented customer-specific and vendor-specific testing of smart cards.
- They can outsource security testing and evaluation to competent, accredited testing laboratories whose expertise can be used by all financial institutions.

The alliance produced a CC profile—a so-called *Smart Card Protection Profile*—of standard security functional requirements and assurance requirements to be addressed by smart cards used in conjunction with sensitive applications such as those encountered in banking industry payment systems. This standard requirements profile was validated and certified by the *Canadian CC Scheme*, and can be used across the financial community.[47] Since then a *UK smart card Protection Profile* for dual-interface authentication cards was developed for uses such as passports.[48]

Consumers can reference such profiles to state their requirements. Vendors can reference such profiles to indicate what they built. Accredited testing laboratories use standard methods to test individual products once for the entire consumer community (not once per consumer institution). The standard tests assess a vendor's claim that its product meets the standard requirements profile. With known confidence, a consumer can then purchase any assessed product that has been independently certified to comply with a specific Protection Profile.

51.4.4.4 Open Source Approach. One approach for software development is use of the Open Source model. According to this approach, software is made publicly available for inspection, for modification of flaws and inefficiencies, and for potential upgrading of capabilities and features. In theory, by the continuous and collective—but uncoordinated and seemingly semirandom—efforts of potentially thousands of autonomous software developers and testers, the public review will improve the quality of the software over time.

The downside of the open source approach is that the degree of trustworthiness achieved by the process is unmeasurable. In certain cases it has been observed that remediation of security flaws is pursued aggressively by major software contributors (especially those that have products that depend on the open source software being remediated); whereas, other software contributors are more inclined to expand open source software functionality rather than fixing discovered security flaws.[49]

To help identify security issues in open source code, the U.S. Department of Homeland Security initiated the *Vulnerability Discovery and Remediation Open Source Hardening Project*. In this project, new approaches for finding critical defects in complex software code sets were developed and used to test open source code to isolate defects and root causes.

In a more recent similar effort,[50] the *Eliminate Vulnerable Code* (eVc) Project[51] uses Internet crawler tools, as well as volunteer identification and/or submissions of bad code, to discover examples of, or citations to, vulnerable open source code. Such discoveries are added to a list shared with eVc members and with developers of the discovered flawed software. The eVc Project hopes that eVc members, sponsors, academia, and more might work with the developers of the discovered flawed software to remediate or to remove the flawed, vulnerable software. By eliminating faulty open source code, it will be easier for other open source projects to reuse higher-quality code that does not contain security flaws such as those identified in the *Open Web Application Security Project* (OWASP) *Top 10* flaws list for Web application security flaws[52] or in the SANS *Top 25* flaws list.[53]

However, the trustworthiness of a product is more than just improved code. Although the open source code will have some degree of trustworthiness developed by the open source process, the incorporation of such code into a product still leads to other factors that influence trustworthiness in the product. Product trustworthiness also depends on vendor processes, such as the quality of design, the protection provided to security features during the delivery of a product from the vendor to the consumer, vendor strategies for maintaining or upgrading security in the face of new threats, and so on.

Several vendors have relied on the standard CC testing and evaluation approach to assess the security and assurance features of their Linux software products. A listing of evaluated Linux products that are certified according to the CC approach can be found online.[54]

51.4.4.5 Hacking. De facto assurance of the underlying security in a product can arise from those who actively probe new products for security flaws. Such probing may arise from internally sanctioned security probing or from unsanctioned, unexpected probing by individuals of ill will. Hacking approaches (ethical or otherwise) do not necessarily follow a consistent or comprehensive approach to evaluating the quality of the security functions and services that are implemented. Hence, the level of assurance achieved is unknown and typically very low.

51.4.4.6 Trade Press. Many trade press publications and magazines conduct reviews of products that pertain to security. Products are tested in ad hoc environments and against private, ad hoc, de jure criteria that vary from product to product and magazine to magazine. Such magazines may rely on unaccredited consultants, staff, or private labs to review products. Some reviews may focus on examining quantitative product details other than security, such as performance or throughput of a product. Tests performed often fall short of assessing the real security aspects of a product. Some reviews rate qualitative parameters, such as product *innovativeness*. Because of the potential lack of quantified testing rigor and potential dissimilarity of evaluation metrics, comparisons of trade press reviews from different sources are difficult. Perhaps most important, no evaluations are made of the confidence (assurance) that can be associated with the soundness of the security implementation.

Examples of publications that provide reviews of security products include *SC* (formerly *Secure Computing) Magazine*,[55] *Network Computing*,[56] *Security Management* magazine,[57] *InfoWorld*,[58] and many more. Some of these trade press initiatives are summarized below.

51.4.4.6.1 SC *Magazine.* In its general test process for individual security-relevant products, *SC Magazine* uses about 50 private criteria codeveloped by SC Lab and the Center for Regional and National Security at Eastern Michigan University. Tests of groups of products focus on operational characteristics. Test and evaluation methods are similar to those used for the CC's lowest assurance levels. *SC Magazine* has also established overall ratings values to indicate how well a product has done against the test criteria and whether the product should be considered for purchase.

Examples of the types of products, or groups of products that have been reviewed by *SC Magazine* include: SIEM, network forensic appliances, multifactor authentication, unified threat managers, risk manager, network sentry, identity enforcement platform, application security manager, password manager, Web application firewall, database activity monitors, secure email gateway, and more. More complete lists of tested items are identified at the magazine's Website.[59]

SC has also produced buyers' guides for groups of products, such as Universal Single Sign-On solutions.

SC Magazine also publishes privately created evaluations of products within certain classes such as cloud identity and access control products, unified threat management solutions, next-generation firewalls, and more.

51.4.4.6.2 Other Publications. *ComputerWorld* and *IT Business Edge* obtain privately created evaluations (often from the same evaluation house) of products from various classes of security-relevant products such as endpoint security products. Individual product evaluations are compared to other products in the same class of products. For security-enabled products, test criteria include performance metrics indicating how security-related processes may impact the overall performance of the system being security-enabled.

NetworkWorld obtains privately created buyers' guides for security-related technology such as identity management in clouds, SSL/VPN tools, and more. Evaluation is based on (a) vendor self-declarations to private, evaluation survey questions in part oriented to security functionality, and (b) evaluator hands-on set up of products.

InformationWeek creates buyers' guides for technology such as gateway antimalware products. Evaluation is based on vendor self-declarations in response to private, evaluation survey questions in part oriented to security functionality.

IT Whitepapers obtains privately created comparisons of security-relevant products or services. Recent evaluation comparisons included, for example, (a) security provisions in major mobile platforms and (b) Managed Security Service Providers.

51.4.4.7 Third-Party Commercial De Jure Assessment Approaches. Early on in the security assessment arena, third-party, commercial, security assessment approaches were sometimes conducted by unaccredited testing houses or consulting houses. These third parties provided relatively low-confidence, so-called surface-level testing. Such testing resulted in vendors of security-relevant products receiving either a *brand mark* or an independent third-party white paper evaluation report that could be referenced in vendors' product brochures and advertisements. Such commercial activities began at a time when there needed to be a lower-cost—albeit lower confidence—alternative to expensive, lengthy, economically inappropriate, government evaluations such as the so-called *Orange Book* evaluations (see Section 51.4.1). These commercial assessment activities were also available to support trade press surveys and magazine reviews of products.

Such nonstandard, third-party approaches are still prevalent. There is a certain qualitative amount of risk reduction achievable by relying on such approaches. They are typically based on simple, one-size-fits-all testing that usually provides minimal, cursory checks of some of the implemented security functions. Some of these tests focus on product details other than security, such as performance or throughput. No evaluation is made of the confidence (assurance) that can be associated with the soundness of the security implementation. At best, these are *black box* approaches wherein products may be examined based only on their outputs relative to stimuli. These approaches have no assessment capabilities based on the fundamental design of the product, or of the engineering principles used by the vendor to build the product.

Many vendors, nonetheless, undergo these types of commercial testing because of the pressures from their competitors' products being so tested. Testing costs are reasonable, but such testing provides no inputs (e.g., evaluation reports) to consumers that can be analyzed to differentiate products. More comprehensive products are not

examined for any of their differentiating capabilities. Instead, such check-mark testing programs merely provide a common-denominator assessment floor for products.

Typical vendor reaction to these types of these nonstandard third-party evaluations and branding programs is that they are often not very good and are often distracting nuisances. Vendors also indicate that unlike a more rigorous, standards-based testing paradigm, such as that based on the CC, these nonstandard, third-party programs do not have processes to help improve the quality of the product under test. Unlike CC testing labs, many vendors do not see these nonstandard testing labs as strategic partners looking to improve the product under test.

Examples of these types of product assessment approaches include the *West Coast Labs Check Mark Program* (Section 51.4.4.7.1), the *ICSA Labs Certified Program* (Section 51.4.4.7.2), and several others (Section 51.4.4.7.3).

51.4.4.7.1 Check Mark Program. The Check Mark program is a private testing service provided by West Coast Labs (www.westcoastlabs.org). Although touted to use *standard* testing criteria in a *standard* testing approach, the Check Mark program establishes private criteria and a private testing methodology that are not recognized by legitimate standards bodies such as ISO. West Coast's private criteria and testing approaches apply to certain types of computer-security products, such as antivirus products, firewall products, and VPN products. The criteria are designed to achieve a basic level of protection against a number of common hostile attacks. West Coast Labs tests products against the applicable Check Mark criteria and, if successfully tested, produces a certificate that shows that specific releases of products meet specific Check Mark criteria.

51.4.4.7.2 ICSA Labs Certified. Another well-known, commercial, security-relevant product branding service is the product certification program conducted by ICSA Labs.[60]

The ICSA approach is similar to the West Coast Labs testing approach. Product performance is tested against specified criteria to assess whether the product can resist the types of common threats and risks specified in the testing criteria. Product testing is typically a checklist-oriented approach geared for nonexpert testers. Testing criteria are developed for a number of classes of products, such as firewalls and antivirus (AV) software. While Check Mark uses private testing criteria, ICSA uses so-called *public criteria*. These public criteria are, however, nonstandard like those of Check Mark since they are created outside the recognized national or international standards-development communities. Instead, ICSA's testing criteria are developed via invited participation.

Products that pass ICSA criteria are entitled to display the ICSA brand mark. Products that fail are reported to their vendors with detailed analysis of the criteria they failed.

Unlike the West Coast certificate, once products are awarded an ICSA certificate, vendors take on the obligation to self-check and to self-declare continued certification of evolutions of the specific version product that passed ICSA testing. Spot checks by ICSA are used to verify that currently shipping products still can pass the ICSA tests.[61]

51.4.4.7.3 Other Third-Party Assessments and Assessors. There are several other organizations that use private, informal methods to create third-party assessments of certain types of security products or security services. Examples follow.

- Gartner issues Magic Quadrant analyses that evaluate how well a product compares to other products in the same class of products and where a product is

positioned in the Magic Quadrant of characteristics. Examples of Magic Quadrant analyses exist for product classes such as Unified Threat Management.

- Forrester issues evaluations of products as well as services in the same class of services. For example, a 15-criteria evaluation of 10 Managed Security Service Providers was conducted using private evaluation criteria.
- NSS Labs uses its own testing laboratory, rather than surveys and questionnaires, to produce best-practices guides, product reviews, comparative test reports, buyer's guides, and more. Such materials pertain to security or security-enabled devices such as, for example, Web browsers, Web application firewalls, next-generation firewalls, endpoint solutions, breach detection systems, intrusion prevention systems, antivirus products, and more. Test criteria are private and tailored to the evaluation to be performed. They may include tests that verify a device's stability, performance, and security effectiveness (e.g., for a firewall not blocking legitimate traffic). Tests may determine the accuracy of a device's security coverage, as well as the usability of the device. NSS Labs publishes its test methodologies for specific devices as well as its own research about security effectiveness. It also publishes vulnerability threat reports. NSS research materials are primarily available to subscription clients.
- Veracode is another third-party assessor available for consumers who are unable to operate a meaningful program for testing the security aspects of vendor products. To support consumer security testing needs, Veracode has established the Veracode Vendor Application Security Testing (VAST) program. In conjunction with a cloud-based testing platform, Veracode provides analysis of security worthiness of vendor application software and its compliance with the customer's security policies and requirements for the application software they buy. It also provides de jure, standards-based, assurance level score requirements that dictate the minimal set of testing that is necessary to be performed. Like the CC paradigm, Veracode looks to establish a partnering relationship with vendors. It discloses its testing results to vendors and feeds remediation guidance back to the software developer. An excellent description of the security evaluation processes to be performed in conjunction with the VAST program is available.[62]
- The AV-Test Institute performs analyses of security products or constituent components.[63] It also performs comparative tests or tests of beta software. Evaluations use AV-Test's own daily-expanding malicious and safe test data samples, and its own analysis tools, as well as private test criteria. The criteria are based on protection (including, e.g., protection against zero-day malware attacks), repair (e.g., in terms of performance in detecting and removing hidden malware), and usability (including, e.g., performance slow-downs of computers using tested products). Based on a scoring system that depends on test results, the Institute awards certification seals of approval. Recertification occurs periodically. Comparative test results for different product classes often appear in the trade press. For example, test results for several Windows 8 virus scanners were recently summarized in the trade press.[64]

51.4.5 Security Assessment Standards for Products. In contrast to the informal, nonstandard product assessment approaches just discussed, formal standards exist for assessing various aspects of security associated with products or services. As detailed earlier, standards exist to assess the overall quality and soundness of product builders (Section 51.4.3.1) and their organization (Section 51.4.3.2) and to stipulate

sound development methods for product builders (Section 51.4.3.3). A determination of security soundness based on these standards typically yield generalized conclusions such as "good" vendors build "good" products.

As detailed in the following sections, other standards can be followed to assess product quality, service quality, or a combination of both the quality of product builders as well as the quality of their products. The latter standards can be used to quantify how well "good" vendors build "good" products (with identifiable and demonstrable assurance levels), how much "better" specific vendors can build even "better" products (with identifiable and demonstrable, generally higher assurance levels), and how comprehensive are the security functionalities within the specific classes of products these vendors build. Similarly, at least one standard exists for verifying that "good" service providers offer "good" services.

Examples of product assessment and service assessment approaches that rely on formal standards are described in Section 51.5 below.

51.5 STANDARD PRODUCT AND SERVICES ASSESSMENT APPROACHES

51.5.1 Government Standard Cryptographic Validation Programs.

The evolving FIPS 140 series of standards[65] specify requirements that are to be satisfied by cryptographic modules. In computer and telecommunications systems used or outsourced by the federal government, the U.S. Government mandates use of validated cryptographic modules for all uses of cryptography that protect "sensitive but unclassified" information. FISMA removes any possibility of receiving a waiver to the government FIPS 140 mandate. The Canadian government requires use of such modules for protection of "designated information." Other nongovernment organizations also look to use validated commercially available, crypto modules.

The FIPS 140 standards provide the basis for the long-standing, popular Cryptographic Module Validation Program (CMVP).[66] This program is focused on testing, evaluating, validating, and certifying hardware and software cryptographic modules. A crypto module may be an embedded component within a product or application; or, it may be a standalone product in and of itself.

A prerequisite to cryptographic module validation is the validation of the correct and complete implementation of the standard cryptographic algorithms used in cryptographic modules. The FIPS-approved standard symmetric key algorithms for encryption include the Advanced Encryption Standard which is the current algorithm of choice, Triple-DES and Skipjack. The Cryptographic Algorithm Validation Program (CAVP)[67] uses CAVP-designed validation test suites for testing, evaluating, and validating standard cryptographic algorithms using processes and methodologies similar to those described for the CMVP below. (Specific tests are described in companion NIST standards.) Details about the CAVP are not offered herein. Suffice it to indicate that thousands of cryptographic algorithm implementations have received algorithm validations to date.

FIPS 140-1 defines four increasing, qualitative levels of security assurance ranging from rock-bottom requirements stipulated in Level 1 to more, sequentially additive, security requirements in each of the subsequently higher levels. Security requirements address 11 different areas that pertain to the design and implementation of a cryptographic module.

FIPS 140-2 and the emerging FIPS 140-3 apply to more contemporary cryptographic module implementations. They incorporate changes in crypto module technology, and in the case of FIPS 140-3 add new requirements and an additional assurance Level 5.

Special security testing laboratories conduct standard cryptographic module conformance testing against requirements stated in FIPS 140-2 or FIPS 140-1. There are numerous such independent, commercial, third-party, so-called Cryptographic and Security Testing laboratories—each individually accredited by NIST's National Voluntary Laboratory Accreditation Program (NVLAP).

Entities that are looking to have crypto module validations performed must contract with any of these laboratories for the service of testing crypto modules and having such evaluations validated. Against the specific crypto module under test, these labs can perform appropriate, standardized conformance tests from a suite of standard conformance tests[68] tailored to the four FIPS 140 security levels.

After validation that the test results from an accredited testing laboratory correctly and completely demonstrate that the crypto module under test conforms to either of the FIPS 140 standards, a CMVP certificate can be issued by either NIST or the Communications Security Establishment Canada (CSEC). The certificate includes an evaluation score determined from the above testing.

NIST maintains a repository list of cryptographic modules—past and present—that have been tested and validated.[69] It should be noted that more products than appear in the repository of validated crypto modules may embed a validated module that is listed in the repository. NIST also maintains a list of crypto modules that are in the process of being evaluated.[70]

The FIPS 140 testing and validation scheme is not as broad, flexible and far-reaching as the CC scheme (Section 51.5.3) that is also used for product evaluation and certification. Why? Because the FIPS 140 scheme is focused on just one class of entities—cryptographic modules—whereas the CC scheme is applicable to any and every class of products as well as to a wider set of possible assurance requirements on the design, implementation, deployment, and maintenance of products. Furthermore, the FIPS 140 security requirements do not map directly to any specific CC security requirements. As such, when an FIPS 140 certificate is required, a CC certificate cannot be substituted.

51.5.2 FedRAMP. The FedRAMP program is a U.S. Government program that supports federal government agencies' adoption of commercial or government cloud services.[71] Such cloud services may be provided according to any one of several different service delivery models, such as Infrastructure as a Service (IaaS) or Platform as a Service (PaaS). FedRAMP is a different type of U.S. Government security assessment and certification program. It is not oriented specifically to security-relevant products. Instead, it is oriented to evaluating service providers and the systems they employ—in particular, cloud service providers.

FedRAMP's standard processes benefit both cloud service providers and cloud service customers. A single, third-party assessment of a cloud service provider's security claims reduces cloud service provider security evaluation costs by eliminating reassessment for different cloud customers. Successful assessments receive government authorization that allows any or all government agencies to use the authorized cloud service provider's offering.

FedRAMP uses a standards-based and risk-based approach that relies on:

1. **Security Requirements Standards** that require cloud providers to implement security capabilities that address standard FedRAMP security requirements.

2. **Security Assessment Standards** that require approved, independent, third parties to audit and to assess cloud service provider offerings and to develop a security assessment report.
3. **Government Authorization Standards** that are used by the FedRAMP Joint Authorization Board to review the security assessment reports and to achieve and to issue consistent standardized authorizations for government agencies to use approved cloud service provider offerings.
4. **Continuous Monitoring Standards** that require specified, standard "Ongoing Assessment & Authorization" activities that include ongoing, periodic reassessment of cloud service provider systems to reaffirm compliance with FedRAMP security requirements, to reaffirm that required security controls are still effective, and, therefore, to retain ongoing authorization.
5. **Authorized Services Lists** that identify cloud service providers that have achieved either an "Authority to Operate" certificate from the authorizing government agency or a "Provisional Authority to Operate" certificate from the FedRAMP Joint Authorization Board and therefore are available for use by any government agency. (Lists of authorized cloud service providers are available at the GSA's Website).[72]

FedRAMP security requirements to be addressed by cloud providers are compliant with FISMA. They are based on security controls enumerated in NITS Special Publication 800-53, Revision 3, which adds cloud-oriented security controls not found in earlier versions of this standard. FedRAMP security assessments add additional controls to FedRAMP assessments that are not included in FISMA assessments.

FedRAMP assessments are initiated either by individual cloud service providers or by government agencies that wish to use cloud services.

A cloud service provider (or the government agency interested in obtaining cloud services) hires a FedRAMP-accredited, commercial, third-party security assessment organization, a so-called Third Party Assessment Organization (3PAO). 3PAOs use FedRAMP standards:

(a) To develop a security assessment plan that is specific to the cloud service provider to be assessed and that specifies the security controls to be assessed, how the controls are implemented within the cloud service provider's environment, security and privacy behavior rules, the boundaries of the cloud service provider's environment, and other security-related plans,

(b) To conduct a security assessment of the cloud service provider's security controls to determine the effectiveness of the implementations of all security controls identified in the security assessment plan, and

(c) To develop a security assessment report to be submitted for authorization approval of the cloud service provider being assessed.

If necessary, the cloud service provider will develop and execute a plan for correcting any weaknesses or flaws identified during the 3PAO security assessment.

If authorization is awarded, then according to FedRAMP's continuous monitoring standard, a 3PAO may conduct annual, follow-up, security reassessments of the cloud service provider's systems in order for the cloud service provider to maintain FedRAMP authorization to continue to be used by government agencies as an approved cloud service provider.

Parties wishing to become FedRAMP 3PAOs must be rigorously evaluated according to FedRAMP standard quality and acceptability criteria used to ensure 3PAO security assessment competency as well as conformity among 3PAOs' security assessment processes.

To leverage the benefits of using an assessed and authorized cloud service provider, government agencies are being encouraged to require, contractually, that their prospective cloud service provider meets FedRAMP security requirements.

Details about FedRAMP, its processes and guidance on how to use the FedRAMP processes are available.[73]

In an effort to be watched, the Cloud Security Alliance is contemplating though its Open Certification Framework (OCF) and Security, Trust & Assurance Registry (STAR) efforts (Section 51.4.4.3.2) to develop what appears could be an international version of a FedRAMP style of cloud service provider assessment and certification.

51.5.3 Common Criteria. Out of the experiences gleaned from the *Orange Book* and other national criteria (Section 51.4.1), a new, commercially driven, standard strategy emerged for testing products and for demonstrating confidence that their security features behave properly. This best-of-all-previous-breeds strategy is based on an international standard, ISO Standard 15408, the *Common Criteria for Information Technology Security Evaluation*, referred to colloquially as the *Common Criteria* (CC).[74]

The CC strategy offers a single, internationally recognized, approach for specification of consumer security requirements and for commercial evaluation of IT product security. Products can be built in any country, evaluated in any other country, and bought with confidence in yet any other country. This CC approach is useful within several international communities. In the United States, it is embraced by the DoD market sector and other communities such as the financial sector.

The CC paradigm presents a standard strategy that overcomes shortcomings of other approaches. It provides a standard way for stipulating (1) the risks of concern, (2) the security functional requirements that must be met in order to mitigate stated risks, and (3) the security assurance requirements that must be met to provide confidence that products are built with desired quality. It provides a recognized, reliable, maintained mechanism to develop trust that:

1. Security requirements are specified correctly,
2. Vendors do not misunderstand the requirements, and
3. Vendors design and manufacture products that address the requirements and provide risk integrity.

The CC provides a flexible process for specifying and testing security requirements for any and all classes, and specific instances, of all existing or future IT products.

In the CC paradigm, consumers' risks and protection requirements are stipulated so that product builders' security solutions can be tested (in a standard way) so as to be able to verify (in a standard way) product compliance with stipulated standard security requirements. Testing is performed by any one of the testing laboratories that are accredited (in a standard way).

Unlike other approaches, the CC provides a way to specify assurance requirements and to evaluate how well they are met. Assurance requirements are extremely important and are typically not considered in other product assessment approaches. Assurance

requirements are the essential ingredients in establishing confidence in implementations of security and in providing the level of trust necessary for economies and governments to rely on new e-business and e-government models.

Trust in products is established by gaining confidence that the security functionalities claimed to be implemented to address specific security functional requirements (1) are effective in satisfying specified security objectives and (2) are implemented correctly and completely and operate as expected. Through security assurance requirements, trust is further established by ensuring that product developers have sound processes and take specified actions to ensure specified quality throughout the life cycle of the product they build, test, deliver, and maintain.

The CC also establishes a standard method to develop common tests and evaluation methods and to use them to verify the security aspects of products via competent, accredited laboratories. CC-based testing helps manufacturers reduce costs by providing a pool of private, licensed, competitive, security testing labs that have consistent, accredited testing quality and competence. Because testing is standardized and recognized worldwide, testing costs are, in effect, capped because product vendors are spared costly country-specific or customer-specific retesting campaigns.

Assessments of products are composed of both analysis and testing of the product. Use of standard evaluation criteria and standard evaluation methodology leads to repeatable, objective test and evaluation results. To remedy assessment-discovered shortcomings and flaws in products, CC testing labs feed back testing results to product vendors and work collaboratively to resolve discovered flaws and shortcomings.

Independent review and validation of CC-based testing and evaluation by accredited, country-specific, national CC validation bodies boosts consumers' confidence even further. Such CC validation bodies verify test results and certify that products were successfully assessed according to the standard testing and evaluation performed by accredited CC testing labs. These bodies are also responsible for maintaining national security testing and evaluation infrastructures that include (a) correct, consistent, credible, competent, and commercial application of CC standards and methodologies by accredited CC testing labs within their country, (b) a government oversight body, and (c) a government security evaluation certificate issuing authority.

The CC specification and testing approach is equally applicable to any and all types of security-relevant products, such as:

(a) Products that implement security technologies (e.g., crypto boxes, intrusion detection/prevention systems, boundary protection devices and systems, audit tools, access control devices, digital signature products, and much more),

(b) Products that are security-enabled (e.g., messaging systems, database management systems, Web e-commerce packages, telecom switches, industrial control systems, and much more), or

(c) Products that support security (e.g., operating systems, certificate management systems, network management systems, smart cards, and more).

An ever-growing list of thousands of products that have been evaluated according to CC-based standard testing and that have received CC certifications is available online.[75] In short, the CC paradigm grows the pool of better-engineered, more acceptable products.

As certified products are altered or upgraded, vendors of such products can pursue a standard assurance maintenance process to retain the assurance level associated

with initial certifications of such products. Information about this process appears elsewhere.[76]

The CC is today's unified choice for developing trust in products. The use of standards tends to increase the product appeal to various nonrelated consumer constituencies. A product builder's security culture is significantly improved because of the required verification of the builder's meticulous and clear focus on security, as well as the builder's security design, development, testing, and maintenance disciplines. Such culture improvement tends to benefit other products built by the vendor.

Some—especially proponents of the 20 Critical Controls paradigm (Section 51.4.2.1)—seem to feel that the CC approach is too flexible, requires too much expertise to capitalize on its flexibility, and would be better if it specified a small, mandatory set of "common-sense" controls of value in "most" situations.

Reciprocally, well-respected consulting houses have concluded that CC evaluation provides a substantial improvement over the nonstandard or informal testing approaches that typically result in seriously undertested software and that, as a group of inconsistent approaches, provide only "apples-to-oranges" comparisons of tested products. The CC paradigm provides an extra level of due diligence. It improves and differentiates products and allows buyers to compare products objectively. It is accepted by mutual agreement in most of the world's largest IT-building and IT-buying countries.

More details about the CC paradigm are given in Section 51.5.3.1. An overview of the two types of security profiles that can be developed by product consumers and product vendors appears in Section 51.5.3.2. A summary for using the CC to document consumers' security needs for products is given in Section 51.5.3.3. Section 51.5.3.4 provides a summary for using the CC to document the security features built into vendor products. More details about the CC testing and evaluation approach appear in Section 51.5.3.5. More detailed descriptions of the underlying CC standards and how to use them appear elsewhere.[77]

51.5.3.1 Common Criteria Paradigm Overview. The CC paradigm uses international standards, processes, and procedures to specify, to test, to validate, and to certify security-related products.

The CC paradigm is a multipurpose scheme for:

- Stipulating security requirements that can be used in product procurements,
- Specifying companion security solutions in products,
- Testing products according to product-tailored—but standard—criteria and testing methodologies using accredited, third-party, commercial testing laboratories,
- Independently validating test results, and
- Providing certificates to tested and validated products that obviate any need for further product retesting for differing consumer constituencies or in all countries that mutually recognize each other's commercial testing capabilities and testing results.

Consumers use a standard methodology, a standard language, and a catalog of standard security requirements to develop security profiles (see Section 51.5.3.3) tailored to the types of products they want to purchase. The profiles stipulate the security functional needs. They also stipulate the confidence or assurance desired in products as well as in product builders' processes from product design through maintenance.

Product builders use the same standard methodology, language, and catalog to develop different types of security profiles (see Section 51.5.3.4) to define their products' security specifications. They define product security specifications in terms of both security functionality and security assurance about the product and the builders' processes. The builders' specifications show how their products meet stated consumer security functionality and assurance needs. Builders' specifications may also show how their products meet any additional builder-claimed security features that go beyond the consumers' stated needs.

51.5.3.2 Specification Elements of the CC Paradigm. The CC standard defines the language used within, the methodologies used to construct, and structures for two types of security profiles that specify security requirements: protection profiles (PPs) and security targets (STs).

PPs help consumers articulate what risks and vulnerabilities are important, while STs help vendors articulate what risks and vulnerabilities are addressed by their products. PPs help consumers articulate what level of protection and confidence they want in products, while STs help vendors articulate what level of protection and confidence are provided by their products. PPs provide a standard and flexible way to transform consumers' security needs and policies into unambiguous, product-neutral, security requirements for the desired security behavior and quality for *any class of IT product*.

Both security functional requirements and assurance requirements are included in these two types of profiles. Functional requirements define desired security behavior and are extracted from a standard CC catalog of hundreds of such requirements. Assurance requirements provide the basis for establishing trust by putting constraints on how well a product or a profile is built. Assurance requirements are also catalogued in a CC standard as requirements that pertain to, for example, (a) configuration management of the product, (b) design, development, delivery, and operation of the product, (c) maintaining assurance throughout the product's life cycle, and (d) much more.

Assurance requirements are bundled into seven standard predefined packages called *Evaluation Assurance Levels* (EALs). Higher assurance level bundles (e.g., EAL4) contain more rigorous assurance requirements needing more rigorous methodologically design and testing than lower, entry-level bundles (e.g., EAL1) that need only functional testing. EALs are enumerated in a monotonically increasing scale that balances (a) the increasing levels of confidence that can be obtained, with (b) the increasing cost and decreasing feasibility of conducting the testing and evaluation necessary to develop a specific, higher level of confidence. The higher the assurance level, the more aspects of a vendor's development process will need to be examined during product test and evaluation.

At the low end, EAL1 can be used, for example, to support the contention that baseline due care has been exercised with regard to protection of personal information and to establish some minimal degree of confidence in correct operation of a product in an environment where the threats to security are not considered very serious. At the highest extreme, EAL7 requires a formally verified design and extensive formal analysis. EAL7 may be applicable to certain highly specific, perhaps one-of-a-kind products, targeted for extremely high-risk situations or where the high value of the assets being protected justifies the extraordinary costs of an evaluation to this level of confidence. Typical commercial products fall in the range from EAL1 to EAL4.

When specifying assurance requirements, the desired level of assurance—and the cost to attain such assurance—should be balanced against factors, such as (1) the value

of the resources to be protected, (2) the risk and extent of possible losses, (3) the level of confidence desired, and (4) any reasonably expected cost and delay factors that may accompany the development, and any subsequent testing, evaluation, and validation of a product at a specific level of assurance.

While a PP stipulates generic, *product-neutral*, security functional and assurance requirements for some specific class of IT product that meets the needs of prospective buyers, an ST is a *product-specific* stipulation of the security functional and assurance requirements addressed by a single, specific product along with information as to how the implemented product meets the stated security requirements.

As stipulated in its accompanying ST, a specific product may claim conformance to one or more PPs as well as to additional product-specific security requirements enumerated within its ST.

Given the breadth, depth, and changeability of possible security requirements that can be stipulated within PPs or STs, the CC provides the ability to stipulate the user requirements and product requirements of virtually an unlimited number of existing and yet-to-be-conceived consumer security needs (PPs) and security product solutions (STs).

The purpose of CC-based testing and evaluation of a product is to evaluate and to confirm that the product meets the product-specific requirements and evaluation criteria contained in the product's ST. The STs' assurance requirements impact the requisite depth and breadth of testing and evaluation.

Many products that are successfully tested and evaluated will go further and have their test and evaluation results validated by a formal, recognized, CC national validation body. Products can then be certified by a government certificate issuing authority and added to the publicly viewable repository that lists CC-certified products.

51.5.3.3 Constructing Protection Profiles. PPs enumerate consumers' security functional and assurance needs that are appropriate and valuable for a specific type of product to mitigate specific risks in a specific threat environment. Each PP states the security problem that a PP-compliant product is intended to solve. It stipulates the security functional requirements that are known to be useful and effective in meeting specific security objectives. It also stipulates the security assurance requirements that provide the confidence desired that products are built with desired quality.

The main contents of a PP include statements about the:

1. **Threats and Vulnerabilities** to which a product will be exposed,
2. **Security Environment** within which a product is to reside, including (a) descriptions of both the IT and non-IT aspects of the environment, (b) assumptions about the product's usage, administration, and management, and (c) all policies, laws, regulations, and rules to which the product must comply within the intended environment,
3. **Security Objectives** to be met either by a product or the environment within which the product operates, including descriptions of which threats and policies are to be addressed by the product and which are to be addressed by the environment,
4. **Security Requirements**, both functional and assurance requirements, selected and refined from the standard CC requirements catalogs, to meet each security objective and assumption, indicating which requirements are addressed by a product and which by the environment within which it operates, and

5. **Rationale** that is provided to justify all decisions and choices made in developing the content within a PP. Rationale includes statements such as how selected security requirements are suitable to counter the enumerated threats, how they comply with enumerated policies and assumptions, and how they map back to the original, underlying, driving needs and threats.

Guidance for developing PPs is available.[78] A list of PPs that were formally validated for correctness is maintained at the CC Website.[79]

51.5.3.4 Constructing Security Target Profiles. Product developers create STs to document detailed information about the security aspects of products they build and to specify how implemented security functions and assurance measures meet consumers' needs.

A ST provides the basis and the evaluation criteria against which testing and evaluation of the product are performed. It also specifies the configuration in which a product is to be tested.

The contents of an ST are similar to a PP. The ST describes the environment within which the product described via an ST is intended to operate. The ST enumerates the (a) threats to the product, (b) policies, laws, and regulations with which the product is claimed to conform, and (c) assumptions about the security aspects of the IT and non-IT environment within which the product is intended to be used. STs delineate the security objectives that the builder of a specific product claims are addressed by the product. STs also enumerate the security requirements that the product builder claims are addressed by the product as well as those requirements to be addressed by the environment within which the product is intended to operate. Some of these requirements may be refinements of generic or product-neutral requirements stipulated in a PP to which the ST claims compliance; and some may be additional security requirements over and above those stipulated in any PPs to which the product claims to be conformant.

However, the ST goes beyond a PP because it also specifies the security functions offered by the product to meet each of the stated security requirements in the ST. It also specifies the assurance measures taken by the product builder to meet all the stated assurance requirements in the ST.

Rationale is also provided for all decisions and choices made in developing the ST content. Rationale justifies all claims made in an ST about the PPs with which the product conforms. In essence, the rationale provided demonstrates that the ST contains an effective and suitable set of countermeasures and that this set of countermeasures is consistent, complete, and cohesive.

By comparing the STs of different products, consumers are better able to compare the security features of competing products. Consumers are able to understand what types of tests and evaluations that a specific product underwent. They also are able to determine whether the configuration in which a product was tested is consistent with the environment into which the product will be deployed.

51.5.3.5 Common Testing of Products. Products characterized by associated product-specific STs can undergo standard CC testing and evaluation using standard testing methodology by any officially accredited/licensed CC testing and evaluation laboratory anywhere in the world. Testing assures that products are correct, complete, well built, and compliant to their security specifications. Testing is conducted under conditions appropriate to the intended normal operation of the product so as to

verify that the implemented security is correct, complete, and compliant under normal operations. Since CC-based testing can demonstrate the traceability of the security aspects of a product back to product user requirements as well as to applicable policies, laws, and regulations, such testing can reduce product users' exposure to potential penalties for noncompliance to security-relevant laws or regulations.

Testing labs and their staff are accredited for their security assessment competence, as specified in a security-testing-laboratory extension[80,81] to an ISO standard[82] specifying general testing laboratory competency requirements. As defined in the CC's *Common Evaluation Methodology* (CEM)[83] standard, evaluators use a standard methodology with structured, formal assessment processes and evaluation actions to carry out a series of standard testing and evaluation activities.

Since testing is done according to standards, consumers get a strong sense that testing is objective and not slanted to benefit the product that was tested. With the additional trust created by CC testing, customers reduce their product acquisition costs by minimizing acceptance testing that duplicates testing already performed on a product by an accredited CC testing lab. Studies indicate that CC testing improves products by eliminating exploitable security flaws and adding or extending necessary security features.

The CEM standard provides a common base for independent, autonomous CC testing laboratories to assess CC profiles and vendor products in the same ways, regardless of the CC testing lab being used. Use of the CEM by all CC testing labs provides a common floor of confidence in similar products that may have been assessed by different CC testing labs. Indeed, the *Common Criteria Arrangement on the Mutual Recognition of the Common Criteria Certifications in the Field of Information Technology Security*, often called just the *Mutual Recognition Arrangement* (MRA), is the multicountry, treaty-level declaration that different countries will recognize the CEM/CC-based assessments that may be conducted by CC testing labs in each other's jurisdictions.

Because the CEM is internationally recognized, the use of customer-unique or country-unique assessment is minimized, if not avoided all together. Assessment costs are minimized, since vendors need only prepare for one testing campaign rather than for a battery of different testing campaigns conducted against different customer-specific, consulting-house-specific, or country-specific, assessment criteria.

Although some consumers and vendors believe that the CC paradigm is complicated, inefficient, or costly, the many who do build and buy validated products see the CC paradigm as the best, proactive way to improve the security and assurance of products bought and sold in the marketplace.

51.6 NOTES

1. "Federal Information Security Management Act of 2002" (Title III of the E-Government Act of 2002), U.S. Public Law 107-347, Section III, December 2002, http://csrc.nist.gov/drivers/documents/FISMA-final.pdf
2. In this chapter, de facto means *accepted as a matter of common practice*; de jure is used idiosyncratically to mean *accepted under written standards* (but not the usual meaning of *according to law*).
3. U.S. Department of Defense Instruction, "DoD Information Technology Security Certification and Accreditation Process (DITSCAP)," 5200.40, December 30, 1997, http://iase.disa.mil/ditscap/i520040.pdf (URL inactive).
4. www.diacap.org

5. PCI Security Standards Council, PCI Standards & Documents, "PCI SSC Data Security Standards Overview," see link for PCI Data Security Standard: https://www.pcisecuritystandards.org/security_standards/
6. The terms *testing*, *evaluation*, and *assessment* are a source of ambiguity and discrepancy in the community. In this chapter, these terms are used interchangeably.
7. ISO is not an acronym; it is a reference to the Greek prefix isos (equal) and ensures that the abbreviation for the organization is identical in all languages. See "Our name" in the page www.iso.org/iso/home/about.htm
8. U.S. Department of Defense Directive 8500.1, "Information Assurance (IA)," October 24, 2002, www.acq.osd.mil/ie/bei/pm/ref-library/dodd/d85001p.pdf
9. Department of Defense Instruction 8500.2, "Information Assurance (IA) Implementation," February 6, 2003, www.dtic.mil/whs/directives/corres/pdf/850002p.pdf
10. www.dtic.mil/whs/directives/corres/pdf/857001m.pdf
11. https://cloudsecurityalliance.org
12. https://www.pcisecuritystandards.org/
13. www.safecode.org/index.php
14. http://omg.org
15. www.webappsec.org
16. www.tmforum.org
17. www.sans.org
18. www.sans.org/critical-security-controls
19. www.microsoft.com/en-us/download/details.aspx?id=18015
20. www.microsoft.com/security/sdl/default.aspx
21. U.S. Department of Defense Instruction, "DoD Information Technology Security Certification and Accreditation Process (DITSCAP)," 5200.40, December 30, 1997, http://iase.disa.mil/ditscap/i520040.pdf (URL inactive).
22. National Security Telecommunications and Information Systems Security Instruction (NSTISSI) No. 1000, "National Information Assurance Certification and Accreditation Process (NIACAP)" (April 2000), www.cnss.gov/Assets/pdf/nstissi_1000.pdf
23. NIST Special Publication 800-37, "Guide for the Security Certification and Accreditation of Federal Information Systems (May 2004), http://csrc.nist.gov/publications/nistpubs/800-37/SP800-37-final.pdf
24. "Information Technology—Security Techniques—Code of Practice for Information Security Management," ISO/IEC 17799 (2005), www.iso.org/iso/catalogue_detail?csnumber=50297
25. Joint Task Force Transformation Initiative (Ron Ross, JTF Leader), "Security and Privacy Controls for Federal Information Systems and Organizations," NIST Special Publication 800-53, Revision 4, April 2013, http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf
26. U.S. Department of Defense, "Trusted Computer System Evaluation Criteria," (TCSEC or Orange Book), DOD5200.28-STD (December 1985), http://csrc.nist.gov/publications/history/dod85.pdf
27. National Computer Security Center, National Security Agency, "Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria"

(TNI), 9800 Savage Rd., Ft. Meade, MD 20755, July 31, 1987, www.ntis.gov/search/product.aspx?ABBR=ADA255422 See also "Trusted Network Interpretation Environments Guideline," www.fas.org/irp/nsa/rainbow/tg011.htm

28. National Computer Security Center, National Security Agency, "Trusted Database Management System Interpretation of the Trusted Computer System Evaluation Criteria," NCSC-TG-021, 9800 Savage Rd., Ft. Meade, MD 20755, April 1991, www.ntis.gov/search/product.aspx?ABBR=ADA393253
29. Office for Official Publications of the European Communities, "Information Technology Security Evaluation Criteria" (ITSEC), Luxembourg (June 1991), www.ssi.gouv.fr/site_documents/ITSEC/ITSEC-uk.pdf
30. www.sans.org/critical-security-controls/cag4.pdf
31. www.sans.org/critical-security-controls/guidelines.php
32. https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf
33. ISO 9000—Quality Management, www.iso.org/iso/iso_9000
34. L. Constantin, "Security Appliances are Riddled with Serious Vulnerabilities, Researcher Says," *PCWorld*, March 15, 2013, www.pcworld.com/article/2030965/researcher-security-appliances-are-riddled-with-serious-vulnerabilities.html
35. for example, P. Wayner, "17 Security Tips for Developers to Safeguard Code," *CIO Newsletter*, February 4, 2013, www.cio.com/article/728207/17_Security_Tips_for_Developers_to_Safeguard_Code?page=2&taxonomyId=3089
36. Mead, N., and B. Hawthorne, "Transitioning from Software to Software Assurance," IEEE Computer Society, *CN Computing Now News Letter*, October 17, 2012, www.computer.org/portal/web/computingnow/security/content?g=53319&type=article&urlTitle=transitioning-from-software-to-software-assurance or http://tinyurl.com/ld4fq5k
37. Summary guidelines for secure software design in cloud and Web environments are available in D. Radcliff, "Code Surety: Secure by Design," *SC Magazine*, March 1, 2012, www.scmagazine.com//code-surety-secure-by-design/article/228646/4/
38. An ambitious *Eliminate Vulnerable Code* project (http://www.digitsec.com/evc.htm) was initiated to crawl the Internet to find pieces of vulnerable, publically available code and to work remove such vulnerable code from the public domain (see Section 51.4.4.4). It is not clear such an effort will succeed.
39. Microsoft, "Security Development Lifecycle," www.microsoft.com/security/sdl/default.aspx
40. K. J. Higgins, "Secure Software Standard In the Spotlight," *Security Dark Reading Applications News Letter*, May 15, 2013, www.darkreading.com/vulnerability/secure-software-standard-in-the-spotligh/240154983
41. IBM Global Technology Services, Thought Leadership White Paper, "Selecting a Managed Security Services Provider: The 10 Most Important Criteria to Consider," May 2011, http://www-03.ibm.com/innovation/us/engines/assets/Selecting_MSS_provider_whitepaper_SEW03026-USEN-00.pdf
42. TechRepublic, "Antivirus Evaluation Guide" provided by GFI Software, January 2013, www.techrepublic.com/resource-library/whitepapers/antivirus-evaluation-guide/
43. L. Musthaler, "Ixia Models Attacks on Your Systems so You Can Test Your Defenses," *Network World Infrastructure Management Newsletter*, February 28, 2013,

www.networkworld.com/newsletters/techexec/2013/030113bestpractices.html?source=NWWNLE_nlt_it_best_practices_2013-03-04

44. www.vpnc.org
45. www.vpnc.org/testing.html
46. https://cloudsecurityalliance.org/star
47. www.cse-cst.gc.ca/its-sti/services/cc/smartcard-sug-v30-eng.html
48. www.commoncriteriaportal.org/files/ppfiles/UNKT-DO-0002%20v1-0.pdf
49. E. Messmer, "Security of Open-Source Software Again Being Scrutinized," *Network World Security News Letter*, March 13, 2013, www. networkworld.com/news/2013/031313-opensource-security-267636.html?source=NWWNLE_nlt_compliance_2013-03-15
50. K. J. Higgins, "Project Finds, Purges Vulnerable Code Snippets From the NET," *Security Dark Reading Database Security News*, May 23, 2012, www.darkreading.com/applications/project-finds-purges-vulnerable-code-sni/240000936
51. https://evc.digitsec.com
52. www.owasp.org/index.php/Category:OWASP_Top_Ten_Project
53. www.sans.org/top25-software-errors
54. www.commoncriteriaportal.org/products
55. www.scmagazine.com
56. www.networkcomputing.com
57. www.securitymagazine.com/
58. www.infoworld.com/
59. www.scmagazine.com/group-tests/section/108/
60. www.icsalabs.com
61. M. E. Kabay was the first Secretary of the Anti-virus Product Developers' (AVPD) Consortium created in the early 1990s by the NCSA (the National Computer Security Association, later renamed ICSA and then TruSecure). In a personal communication, he reported that anti-virus vendors were using radically different standards for their definition of anti-virus coverage; with a few thousand known viruses in the world, some vendors claimed to combat thousands and others only hundreds. When the NCSA told the new AVPD members that the first test would require them to identify and block only 75 percent of Joe Wells' WildList, many of the technical staff from the companies laughed. They didn't laugh so much when many of their products failed the initial round! Over the next rounds of testing (4 per year), the percentage of coverage of the WildList rose to 100 percent and then went on to include other viruses in the "zoo" of laboratory-identified viruses that had not yet spread to ordinary computer users' computers. According to Kabay, the NCSA program led to major improvements in the quality of anti-virus products and increased clarity. He argues that the insights and methods developed in the AVPD were then applied successfully to other security products.
62. Veracode, "Five Best Practices of Vendor Application Security Management," White Paper, available in "The Secure Software Supply Chain Toolkit," *All Things Security* Blog, September 20, 2012, www.veracode.com/blog/2012/09/the-secure-software-supply-chain-toolkit
63. www.av-test.org

64. "Ms. Smith," "AV-Test issues first Windows 8 antivirus solution ratings," *Network World*, April 8, 2013, www.networkworld.com/community/blog/av-test-issues-first-windows-8-antivirus-solution-ratings?source=NWWNLE_nlt_daily_pm_2013-04-09_testA6
65. NIST, Computer Security Division, Computer Security Resource Center, CMVP, "Standards," http://csrc.nist.gov/groups/STM/cmvp/standards.html
66. NIST, Computer Security Division, Computer Security Resource Center, CMVP, "Cryptographic Module Validation Program," http://csrc.nist.gov/groups/STM/cmvp/ index.html
67. NIST, Computer Security Division, Computer Security Resource Center, CAVP, "Cryptographic Algorithm Validation Program," http://csrc.nist.gov/groups/STM/cavp/ index.html
68. CMVP Program Staff, "Derived Test Requirements for FIPS PUB 140-2, *Security Requirements for Cryptographic Modules*" (draft), published by NIST, January 4, 2011, http://csrc.nist.gov/groups/STM/cmvp/documents/fips140-2/FIPS1402DTR.pdf
69. NIST, Computer Security Division, Computer Security Resource Center, CMVP, "Module Validation Lists," updated periodically as new validations emerge, last updated August 27, 2013, http://csrc.nist.gov/groups/STM/cmvp/validation.html
70. http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140InProcess.pdf
71. www.gsa.gov/portal/category/102371
72. www.gsa.gov/portal/content/131931
73. U.S. DoD, DHS, GSA, NIST, "Guide to Understanding FedRAMP," Version 1.2, April 22, 2013 (and subsequent versions), www.gsa.gov/portal/category/102375 ("Key Documents" link on "About FedRAMP")
74. The Common Criteria Portal, www.commoncriteriaportal.org
75. Common Criteria Certified Products, www.commoncriteriaportal.org/products
76. S. Amiri, "Significance of Common Criteria Assurance Maintenance (AMA)," *SC Magazine* (September 22, 2004). www.scmagazine.com/significance-of-common-criteria-assurance-maintenance-ama/article/31385
77. P. Brusil, "Security Standards for Products," Chapter 51 in *Computer Security Handbook* (5th ed.), ed. S. Bosworth, M. E. Kabay, and E. Whyne (Wiley, 2009).
78. Technical Report ISO/IEC TR 15446, "Information Technology—Security Techniques—Guide for the Production of Protection Profiles and Security Targets" ISO/IEC TR 15446:2004, July 2004, http://standards.iso.org/ittf/PubliclyAvailableStandards/c039690_ISO_IEC_TR_15446_2004(E).zip [Available in the extensive list of "Freely Available Standards" http://isotc.iso.org/livelink/livelink/fetch/2000/2489/Ittf_Home/PubliclyAvailableStandards.htm]
79. Common Criteria Certified Products, www.commoncriteriaportal.org/products
80. ISO/IEC Technical Report 13233, "Information Technology Interpretation of Accreditation Requirements in Guide 25 Accreditation of Information Technology and Telecommunications Testing Laboratories for Software and Protocol Testing Services," available for purchase online, www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=21468
81. NIST Handbook 150-20, "Information Technology Security Testing—Common Criteria," www.nist.gov/nvlap/upload/NIST-HB-150-20-2005-1.pdf—a techno logy-specific extension to J. L. Cigler and V. R. White, eds., NIST Handbook 150,

"National Voluntary Laboratory Accreditation Program—Procedures and General Requirements," U.S. Department of Commerce, Technology Administration, National Institute of Standards and Technology (Washington, DC: US Government Printing Office, 1994), www.nist.gov/nvlap/upload/nist-handbook-150.pdf

82. ISO/IEC/EN Guide 17025, "General Requirements for the Competence of Calibration and Testing Laboratories," 2011, www.fasor.com/iso25
83. Common Evaluation Methodology Editorial Board, "Common Methodology for Information Technology Security Evaluation: Evaluation Methodology," Version 3.1, Revision 3, July 2009, www.commoncriteriaportal.org/files/ccfiles/CEMV3.1R3.pdf

INTRODUCTION TO PART V

DETECTING SECURITY BREACHES

No matter how well we implement security mechanisms, we are facing human opponents who may counter our best efforts until we can respond appropriately. How do security and network administrators find out if there has been a breach of security? How can they evaluate their own defenses before they are penetrated? This part includes chapters on:

52. **Application Controls.** Application-software security and logging
53. **Monitoring and Control Systems.** System logging and data reduction methods
54. **Security Audits.** Measuring compliance with explicit policies and with industry standards
55. **Cyberinvestigation.** Organizing effective digital forensic studies of observed or suspected security breaches, for internal use, and for cooperation with law enforcement

CHAPTER 52

APPLICATION CONTROLS

Myles Walsh and Susan Baumes

52.1 INTRODUCTION **52·1**

52.2 SYSTEMS ARCHITECTURE AND DESIGN **52·2**

52.3 INPUT AND ENTITLEMENTS **52·2**

52.3.1 Input 52·3
52.3.2 User Entitlements 52·4

52.4 DATABASES AND UNDERLYING DATA CONTROLS **52·5**

52.4.1 Data Corruption 52·6
52.4.2 Database Management Subsystems 52·7
52.4.3 Distributed Databases 52·8
52.4.4 Recovery and Restart 52·12

52.5 BATCH FILES **52·13**

52.5.1 Backup File Creation 52·13
52.5.2 Audit Controls 52·13

52.6 DATA INTEGRITY AND VALIDATION **52·13**

52.6.1 Validation Controls 52·13

52.7 CONCLUDING REMARKS **52·16**

52.8 FURTHER READING **52·16**

52.1 INTRODUCTION. Application controls can be broad and multifaceted. Given the complexity of many applications, no single system can provide all of the controls needed. As part of the design of a new or changes to an existing application, there are a number of controls that must be considered to ensure that the underlying data is protected against threats to the six fundamental attributes of information (confidentiality, control, integrity, authenticity, availability, and utility) as discussed in Chapter 3 in this *Handbook*. As with all other information security areas, a defense-in-depth strategy for application controls is a necessity. As an example, should a Web server be publically available? If not, then that control would have to be considered in the design of the application.

The basis of application controls fall into three areas:

1. Systems architecture and design
2. User input and entitlements
3. Database and underlying data controls

There is tremendous complexity in modern systems due to architecture changes and cost pressure on IT departments. Because of this complexity, there is the potential to have so many controls that they end up overriding one another and in some cases

limiting their usefulness. It is easiest to design a net new application and consider all of the controls necessary, but the same due diligence is required for changes to existing applications whether new functionality is being added or there are changes to business processes that may require large entitlement changes. All changes in applications need to follow a strict design method whether using agile, waterfall, or some other methodology.

It should be noted that every architecture tier plays a role in applications and each tier must have some controls in place.

Although this chapter focuses on data-processing applications, similar issues arise in supervisory control and data acquisition (SCADA) systems for real-time process control.

52.2 SYSTEMS ARCHITECTURE AND DESIGN.

Architecture and design provides the first line of defense for applications. The configuration of the design and systems are paramount. Protection of and how the controls effect an application must be considered in the underlying architecture and design.

Architecture controls fall into a number of categories, but they encompass all of the tiers of a system supporting an application. These are standard best practices for systems.

Things to consider when designing system controls for applications include:

1. What underlying processes need to run on the system in support of the application? **Recommendation**: Limit processes to only those that are required.
2. What ports need to be open in order for the system to run correctly in support of the application? **Recommendation**: Limit ports to only those that are required. See Chapter 26 in this *Handbook* for further discussion of gateway security.
3. Who has direct access to the system and why? **Recommendations**: Practice least privilege and require special approval for higher risk or administrative access. Regularly review all access, but if access is higher risk, more frequent review and validation of access may be required.
4. What controls and processes will be in place for support of the application? Similarly, what controls will be in place for change control or patching? **Recommendations**: Define, follow, and validate change control processes for the systems and applications that those systems require. See Chapter 40 in this *Handbook* for more details of patch management.
5. How are the interacting systems validated with each other to ensure they are who they say they are? **Recommendations**: Digital certificates and two-way SSL can be used to ensure the validity of the underlying systems. Know and understand what systems can communicate with other systems and why.

52.3 INPUT AND ENTITLEMENTS.

From an application development perspective, the biggest control that must be considered is user input. User input falls into two distinct areas. What rights or entitlements does or should the user have and what are the limits for the input that the user is allowed to have? Business requirements of the applications generally drive these specifications. An example could be the input field of an address. The business may stipulate that the field be 50 characters and that only certain types of users be granted access to input, change, or delete the field.

Questions to ask regarding entitlements include:

- What is a user or groups of users allowed to do in the application?
- What should the user be able to view, create, update, or delete?
- Are there specific functionalities that a particular user or groups of user are not allowed to perform because of their current role within an organization?
- From a reporting perspective, the same questions need to be asked.

As part of application design all of those questions need to be answered based on the business requirements.

On the other hand, each input field within those entitlements must be defined within certain limits and those limits must be carried from the initial input to the underlying data structure. Both the application and the underlying database can control the limits and data entered.

Continuing with the above example, the 50-character address field should only accept 50 characters. The field length should be 50 and there must be input validation completed to ensure only valid or approved characters are entered. This control can be based both at the field level of the application as well as the database where the data will be stored. Limiting the length of inputs is a key defense against injection attacks and buffer overflows.

Typically, data validation has occurred closer to user input simply due to speed of networks, and more recently, real-time user validation provides a better experience for the user (e.g., validation as the field is being entered versus at the "submit" button of a Web-based form). Best practice requires server-based validation of input prior to processing. Initial checks can be completed on a Web form or field within a desktop application, but server-side validation is better since it is closer to internal systems.

Systems can accept input from both users and systems; for example, a batch process may take input from one source, manipulate it, and carry it to another source. Input, whether by a user or system, must be validated.

Database controls work in support of application controls. If, for example, field-level validation were not occurring in an application, the underlying database could be used to protect against bad data. As with the example above, if the field-level validation and the server-side validation failed, the database field can be defined as a 50-character field. Most database systems will truncate and ignore additional data that is entered into a field.

52.3.1 Input. As stated above, all input must be cleansed and validated to ensure that only acceptable values are received and processed. This is true for user input as well as systems input. The same thought must be used for a database perspective. Obviously, there are challenges with this method. If all 50-character fields for addresses must be updated to accept 60, multiple changes and more thorough testing are required.

Organizations may decide that the applications data is not critical and design it with looser controls. This can lead to issues simply because an attacker may use an application with fewer controls as a jumping point to get to an application with more sensitive data.

One cannot stress enough the importance of data validation for all user input. Many attacks on systems could have been prevented if proper and comprehensive data validation were completed (e.g., many cross-site scripting (XSS) and cross-site request forgery (XSRF) attacks [see Chapter 16 in this *Handbook*] can be prevented

with validation). Not accepting certain known malicious characters, encoding certain characters, blocking certain string combinations, and understanding application attack vectors are important to keep applications secure and maintain integrity. Application designers need to work with their business partners and explain the issues with allowing certain types and sequences of characters.

As with all information security controls, multiple levels of control are needed. It is not appropriate to build an application where there is only field-level validation of input. Multiple controls need to be used in order to provide a defense-in-depth strategy for applications.

52.3.2 User Entitlements. Especially daunting are user entitlements for applications. Defining and limiting access while enabling users to complete their required tasks can be a formidable challenge. Application controls are paramount when considering what a user can do when within an application. Entitlements are also important when considering segregation of duties within the same application, as users change roles within an organization, or even what, if any, access a developer would have to a production environment.

User access rights are the entitlements granted to users or groups to perform functions within an application. As an example, a user that is supposed to have read only access to data should not be granted rights to an account that allows updates. Good design uses systems, application, and database controls to allow or restrict users to perform different functions.

If an individual's role within an organization is as a developer, should they be allowed to migrate code to production? Best practices tell us that they should not; however, every organization is different and development teams may also migrate code to production. The important thing is to:

1. Understand where entitlements are granted, to whom, and why.
2. Verify entitlements on a cyclical basis.
3. From a development perspective, have entitlements listed as requirements, in the test plan, and validate that the tests were successful.

52.3.2.1 Practice Least Privilege. As a best practice, apply least privilege to user entitlements, as it should be applied in all other areas of technology in order to maintain control over data and processes. How the underlying system is architected, the size of the organization, and regulatory controls will dictate the feasibility of achieving the best-practice state.

In many instances, business requirements do not consider which specific types of users should have access to which functionality. This lack is not sustainable, simply because applications are constantly changing, functionality being added, legacy architecture being updated, and new technologies being included.

Careful consideration must be given to grant users only the permissions required to do their job function(s) and controls. Comprehensive testing must also be conducted to ensure access creep does not occur as changes are made to the application.

52.3.2.2 Solving the Problem. For smaller applications, one can easily define the different business requirements with the functions that should have access. A simple matrix may do, for example:

Business Function	Business Role	Technology	Application or database role
Add new account	Phone representative	Insert into certain tables	App: phone rep DB: phone rep
Update account	Phone rep (with review) or data input associate	Update certain tables	App: update account DB: update account
Confirm phone representative	Phone rep manager	Read certain tables Update certain tables	App: phone rep manager DB: phone rep manager

Once the business requirements are defined and put into a matrix, it becomes much easier to manage existing functions and any future changes and to understand where the critical functions are to grant only the required privileges.

52.4 DATABASES AND UNDERLYING DATA CONTROLS. Database technology allows for granular control of information though the use of role-based access controls (RBAC), as discussed in Chapter 9 in this *Handbook*. In addition, databases provide a multifaceted approach to application control. Not only can all users be segregated based on business need, but the database design can provide a level of control using table design and layout. Good database design and use can help protect data security. Relational databases are based on an underlying theoretical model by E. F. Codd, developed in the 1970s, and widely used today and for the foreseeable future.

Relational database management systems (RDBMS) are sets of programs that provide database administrators (DBAs) with the tools to perform the following tasks:

- Create database structures (file or table layouts, and screens or forms).
- Enter information into the structures.
- Establish cross-references among the files or tables.
- Manipulate (sort, index, and summarize) the information in the structures.
- Import information from nonrelational database structures and export information to nondatabase structures. This allows for interfaces between applications using an RDBMS, and applications using conventional files structures
- Data-repository and data-dictionary tools. These tools document databases and may include descriptive information (*metadata*) about all of the resources included in an information-systems environment. Data dictionaries can support documentation standards. Common file or table structures, common program modules, and common field definitions contribute significantly to the reduction of confusion, and of communication breakdowns in an enterprise
- Databases move much of the dynamically changing requirements for data formats and data edits out of the source code of application programs and into easily localized and edited single criteria. Instead of having programmers hunt down and alter multiple occurrences of data restrictions, DBAs can collaborate with the programmers to change individual criteria with the assurance that all programs using the RDBMS will successfully execute the changes.

All of these tools help maintain the security and the integrity of the database.

There are two primary methods of organizing files and databases for access and processing: *batch* and *online*. Batch files, which can be written in various *job-control languages* (JCLs), carry out instructions without further human intervention once they are set in execution, whereas online methods involve human interaction with the programs at multiple points. Protection of online files and databases requires additional planning when the systems using them are being architected and special precautions when they are being used. Protection of batch files is more straightforward because the creation of backup copies is an inherent part of the batch process.

As previously discussed, data validation prior to entering information into a database must be completed regardless if the entry is via user or system. Every attempt should be made to assure that data are entered or transferred and validated expeditiously. Failure to get correct data into a system quickly delays the creation and display of the information, the very purpose of almost all systems. When improper attention is directed to validation of the data as it enters the system, incorrect information is created by the system. The acronym *GIGO* (*garbage in, garbage out*) has long become part of the national lexicon of everyday language. Improper validation and allowing of garbage into a database may have downstream impact to other applications or processes that use the data. Validation at all tiers of an application is critical.

For related topics, see Chapter 38 in this *Handbook* on writing secure code, Chapter 39 on software development and quality assurance, Chapter 40 on managing software patches, and Chapter 47 on operations security and production controls.

52.4.1 Data Corruption. Data corruption implies incorrect data—a violation of the *integrity* of the data. Corruption can occur because of physical factors or logical errors in programs. Good application controls dictate that databases and their underlying architecture be protected from data corruption. Data can be damaged through hardware problems, errors in software, or during concurrent access by two or more users or processes.

52.4.1.1 Physical Data Corruption. Physical corruption occurs through breakdown or other failures of hardware such as computers and network equipment, especially of mass-storage devices such as magnetic disks, tapes and cartridges, or optical drives. Data corruption during transmission can occur through electromagnetic perturbations of communications cables or radio-frequency noise that affects wireless transmission. Fiber-optic cables are susceptible to cross-talk and to disruption caused by physical bends or kinks in the cables; optical disks are susceptible to dirt and abrasion; magnetic disk drives can be harmed by physical jolts (dropping, intense vibration); solid-state drives do not do well when immersed in liquid.

In addition to problems in the transmission or storage media, improper settings or the effects of wear in the equipment can cause errors. Examples include misalignment of magnetic heads on disks and tapes; bad contacts in wireless transmission equipment, causing noise; and improper positioning of lasers in optical media. See Chapters 22 and 23 for extensive discussion of physical factors that can cause data corruption.

Physical corruption characteristically shows data written inappropriately across blocks, rather than in single fields defined by software applications. Such physical corruption usually has no relationship with the damaged records other than physical proximity on the disk; therefore, a cardinal sign of physical corruption is damage to files from completely different applications—a block of database records contiguous to a block of text files, for example.

52.4.1.2 Logical Data Corruption. Logical corruption occurs through programming errors such as incorrect sums, bad arithmetic formulas, incorrect logical conditions, bad data in look-up tables, out-of-bounds conditions allowing reading and writing in the wrong areas of memory, lack of data validation, and the effects of malicious software. See Chapter 38 in this *Handbook* for a discussion of many types of such errors.

Logical corruption characteristically shows errors in the same field in multiple records. Another class of logical errors shows bad values for boundary conditions (e.g., the smallest or largest possible value) but correct data within the range. Such errors rarely cross application boundaries unless there is a logical relationship among the damaged files; for example, an error in a spreadsheet (almost always due to input or programming errors by the user) may propagate to text files if object linking and embedding are used to insert the spreadsheet results into the document. Other documents will remain unaffected by such logical errors.

For more extensive discussion of preventing, locating, and correcting logical errors, see Chapter 39 on software development and quality assurance.

52.4.2 Database Management Subsystems. In the 1960s, complex systems, such as those based on accounting records, forced programmers to define their own file structures to represent relationships among the data; for example, an order header file would be linked to the corresponding order detail records through hard-coded relationships in the application programs. Each programmer or programming team had to define its own data access paths and code them explicitly. Coordinating access to multiple individually named files caused headaches for programmers. For example, it was easy to make errors such as forgetting to increment counters representing how many detail records corresponded to a master (index) record (e.g., line counts in order headers disagreeing with the actual number of detail lines). Because there was no particular protection for the files in such systems, it was easy to replace or delete individual files by mistake, leading to massive logical corruption. Deletion of header records could leave a fragment of inaccessible details (*orphans*) in associated files. Programs had to keep pointers up-to-date for forward and backward chaining. Backups sometimes failed through operator error so that not all related files were included. Furthermore, every system had its own unique methods for managing the data, causing maintenance and learning headaches.

The late 1960s saw many large programming shops defining their data access methods by using library routines that all programmers could share, but there was still a heavy investment in learning the new rules whenever a programmer changed jobs. Following the publication of E. F. Codd's model for relational databases in the early 1970s, the programming field saw explosive growth in implementation of database management subsystems, where the interface to the database controlled utility functions such as indexing, pointing, and chaining.

With proper configuration a typical DBMS protects all files, called *datasets,* against accidental erasure, and would force all reads, writes, appends, and locks to be mediated by the DBMS routines. General users should not have access to a DB's underlying infrastructure—only to the data through the DBMS.

The DBMS provides controlled access to data and also typically provides backup utilities to ensure that all datasets are copied together and to prevent accidental restoration of the wrong version of a dataset. Finally, a DBMS usually provides logging

facilities to keep records of different types of access to the database. All of this functionality provides control to the data.

One of the most important rules enforced by a DBMS is *referential integrity* to prevent common logical data corruption, automatic *uniqueness constraints* to preclude duplicate and conflicting records, and *locking* for safe concurrent access.

- ***Referential Integrity.*** Referential integrity in DBMS design ensures that every dependent record has a primary key value that matches an existing primary key in the master file.
  - In an order database, for example, order numbers in a header file (often called the *order master*) are the primary keys.
  - Each order header record contains unique information about an order, such as customer number, date placed, total price of materials, taxes, and shipping costs.
  - The order detail file contains the dependent records, each of which can be located using its order number as the primary key.
  - Each detail record contains information about a specific part of the corresponding order such as an item number, quantity ordered, price, extended price, and special charges or discounts.
  - If an order header record is to be deleted, all the order detail records must first be deleted; otherwise, the detail records would be left with no way to locate them through their primary key value.
  - Similarly, no detail record can be added unless the master record with the same primary key already exists.
- ***Uniqueness Constraints*.** A modern DBMS allows configuration of nonrepeating primary keys; for example, in an order database, the order number would typically be a nonrepeating or unique key because there should never be two orders with the same identifying number. Setting the uniqueness property would preclude adding a second header record with the same value in the order number field as another order.

52.4.3 Distributed Databases.

A *distributed database* is one that is stored in different databases on different computer platforms. Databases can be distributed over many sites, using several different architectures.

The simplest and most susceptible to failure is the single database server that houses the database and shares its contents among several local area networks. Whenever the server or the database is offline, or during a crash, all clients lose access to the database. The backup and recovery of a single-server database follows the procedures described earlier.

A second architecture is the *replicated database server*, in which there are duplicate copies of the database on multiple servers. This type of system allows clients to access information from any of several copies of the database. With a replicated server environment, clients can still have access to the database, when one of the servers is offline. Such replicated systems usually have an additional benefit of being able to balance transaction traffic so as to keep any one server from being overwhelmed. The trade-off is the database synchronization process that increases the complexity of the replicated server architecture. Although the backup and recovery of a distributed database can also follow the procedures described earlier, it is complicated by the synchronization requirement.

A third architecture is known as a *partitioned database server*, in which specific subsets of a database are on two or more database servers. There could be a marketing database on one server, an accounting and finance database on a second server, and an inventory database on a third server. Because it is generally not possible to have mutually exclusive subsets of databases, synchronization must still be part of this architecture. Backup and recovery of partitioned databases requires application of the procedures described above for each of the subsets.

There is another distributed architecture known as *federated database servers*. This architecture is used in support of databases on two or more servers, made up of ordinarily incompatible storage models such as hierarchical and relational models supported by different DBMSs. With the federated architecture, a single unified database definition, or schema, is created and stored on a combined database server. That server acts as an interface between application programs and the databases residing on the other servers. Queries and other transactions are sent to the combined database server, which translates these into queries and transactions to the underlying databases. Responses then return from the underlying databases to the unified schema to formulate a response to the user. Backup and recovery procedures such as those described above are used for the underlying databases.

Although federated database architecture can be complex and expensive to maintain, it can be less complex and less expensive than the process of supporting application programs for each of the underlying DBMSs. Federated architecture is often used in the implementation of data warehouses, which are used to extract information from multiple internal databases, as well as external databases, to support management decision making.

52.4.3.1 Database Controls to Maintain Integrity. Concurrent access to data can cause problems of integrity unless the access is *serialized* to ensure that one process at a time accesses any given resource.

52.4.3.2 Lock on Update. When more than one user accesses a database, it is possible to experience conflicts over the use of specific records. A classic example occurs in an inventory database, where there are 15 units of part 1 in the inventory.

- User Albert needs to take five units of part 1 out of inventory, leaving a total of 10. The inventory program reads the inventory record for part 1 and modifies record 1 to show only 10 units.
- However, if while this is going on user Betty needs three units of part 1, the inventory record still shows 15 units available because user Albert has not yet updated that record.
- After user Albert's program completes its update, the record shows 10 units available, but after user Betty's program overwrites that record to show 12 available units for part 1, the inventory total is off by five units.
- This situation is an example of a *race condition* as described in Section 39.4.1.6 of Chapter 39 in this *Handbook*.

To avoid this kind of logical corruption, a DBMS provides facilities for *locking* parts of the database. In the inventory example, user A's program would lock the inventory record for part 1 (or the entire inventory dataset) until the update is completed. That

way, user B's program would have to wait for user A's program to *unlock* the data before being able to act on the inventory record.

The obvious symptom of a bad locking strategy is a discrepancy between the database value and the real-world value. However, such a discrepancy is not by itself proof of logical corruption, because the same divergence could arise from events in the real world not reflected in the database. In the inventory example, the actual inventory may have been reduced by theft, or increased by an unrecorded addition of materials.

52.4.3.3 Unconditional versus Conditional Locking. There are two types of locking strategy: *conditional* locking and *unconditional* locking.

- Conditional locking attempts to obtain a lock, but if the required record or the entire dataset is already locked, the DBMS returns control to the calling program with a status indicator of this condition. The application program can then be written to loop until the lock is obtained.
- The unconditional lock request hangs the program until the lock is granted. The DBMS or the operating system provides automatic queuing using a first-in, first-out queue.

52.4.3.4 Deadlocks. Unconditional locking carries risks if multiple resources are locked by programs.

- For example, if program A unconditionally locks resource 1 and program B locks resource 2, trouble will occur when program A then attempts to lock resource 2, while program B tries to lock resource 1.
- Neither program will release the resource it has locked until it is released, and so both will wait forever or until one of the programs is forcibly terminated.
- Such a situation is known as a *deadlock* or more colorfully as a *deadly embrace.*
- If programmers insist on using unconditional locking, the deadlock prevention strategy is to ensure that all programs accessing the database must lock resources in the same order (e.g., lock 1, then lock 2) and must unlock in the reverse order (unlock 2, then unlock 1).

Other strategies are to keep transactions as short as possible, and to avoid the necessity for operator interactions that would keep the records locked for long periods of time.

52.4.3.5 Two-Phase Commit. Sometimes many records or datasets must be locked for complex transactions to be completed. For example, in a hospital's clinical systems database, discharging a patient could require modifications in datasets such as the patient-master, treatment-detail, nursing assignment master and details, doctor assignment master and details, and datasets for the financial functions. Locking everything that might be needed and waiting for a human being to enter all the appropriate data could take seconds to minutes, during which all the affected records would be locked and unavailable to everyone else on the system. In the extreme, if an operator were to leave in the middle of a transaction, other users could be blocked out of large parts of the database for an indeterminate length of time, even for hours. The delays resulting from such locking around human intervention led to the principle that *no transaction can be allowed to lock around a human intervention.*

Another problem with locking around human intervention is that a system failure could terminate processing while the database was in an inconsistent state. For example, in the inventory case, the DBMS might have updated the order detail by adding an item, but not yet updated the order header to show the new total cost of the order. To reduce the likelihood of such an occurrence, a DBMS can support the two-phase commit as an aid to making changes as fast as possible, and thus reducing the window of vulnerability for interruption.

- In the two-phase commit, the DBMS obtains copies of all the records needed when the operator begins a transaction.
- Once the operator has taken all necessary steps for completing the transaction, the DBMS locks and reads all the changed records again and compares the current values with the initial values; if there are no differences, the DBMS makes all the required changes and immediately unlocks all the records.
- However, if the current values have been modified since the operator requested the initial copies, then some other process has been active and so the DBMS reports that verification is needed.
- The operator is typically given the choice of how to proceed; for example, if there are no items left in inventory, the order may have to be delayed or canceled, whereas if there are enough items, the operator need merely reinitiate the transaction with the new initial values.

52.4.3.6 Database Backup Files and System Logs. When information in an online file or database is updated, the old information, that is, the information that was in the record before the change was made, is overlaid; unless steps are taken, it disappears without a trace. For this reason, many DBMSs allow an image of the original record to be copied to the *transaction log file*. In other cases, where the exact history of a particular group of records must be preserved, as with insurance and medical data, an application may append new records but not delete old ones.

If online files and databases were never damaged or lost, data loss would be of no concern with proper database configuration. However, in an imperfect world, steps must be taken so that damaged or lost online files and databases can be recovered. For further details of data backup and recovery, see Chapter 57 in this *Handbook.*

In order to recover online files and databases, it is first necessary to make periodic backup copies, and to make log copies of records that have been updated in the time between making the backups. How often a backup copy is made depends on how dynamic the files or databases are. In most enterprises, a significant number of total or partial files and databases are copied daily. It is not uncommon for computer operations departments to spend several hours each day doing backups. Whenever a backup copy of a file or database is created, there are two correct copies at that point in time. To explain how online file and database backup and system logging work, think of a single online file. The file is taken offline, and is no longer accessible to online transactions at 4 AM, and a copy is made. Both copies of this file are identical at that time. At 6 AM, the backup is complete and the original disk file is put back online. Transactions start to update that file again. From that point on, with each update transaction, the differences between that file and its backup increase.

At 2 PM, if for some reason the file is no longer usable, the backup file is then eight hours behind. At this point, the log file becomes critical in the restoration process. The

log file is in sequence by time and contains copies of the updated records both before and after update was performed. It also contains copies of the transaction record.

52.4.4 Recovery and Restart. After it has been determined that an online file or database has been corrupted or destroyed, a procedure known as recovery and restart is initiated. The first step in this procedure is to copy the backup back to disk to create a new original as of the time of that backup. The next step uses the log file to reapply, in time sequence, all the transactions that had been executed since the backup copy was made.

Contemporary database systems may have files that cannot be taken offline. They are online 24 hours a day, seven days a week. In order to make backup copies, parts of the database are copied periodically (dynamic backup). Conceptually, a database can be broken up into parts. Each part can be backed up separately at different time periods. A number of schemes can be devised to back up some of the records. For example, copy to a backup file every fifth record of the database in one time period, say records 5, 10, 15, 20, etc. A bit later, copy to another backup file records 1, 6, 11, 16, 21, and so on. If a conflict occurs between copying the record for backup and a transaction attempting to update the record, have an appropriate procedure established to let one or the other take place first. Recovery and restart, as well as backout procedures, work the same way, with the additional complexity of establishing priorities for which gets done first when conflicts occur. Even though these conflicts increase complexity, they are resolved when the procedures for recovery, restart, and backout are created.

Another approach to backups of such critical files is the *shadow copy*, which copies data even when files are in use. Shadow copy software bypasses elements of the file system to be able to access locked records.

52.4.4.1 Backout. The log file is also used in a process known as *backout*. This process is initiated when online update transactions fail to complete after making incomplete or partial updates to files or databases. For example, an update transaction in which there are additions to three fields in three separate files is supposed to take place. After making two out of the three updates, the transaction terminates abnormally because of a program malfunction. Eventually, the program gets corrected, but something has to be done to undo the two partial updates. Otherwise, when the program is fixed and the transaction is rerun, those two updates would be reapplied, which would result in erroneous duplication. The backout procedure is initiated for those transactions which, if they had not completed properly, would generate errors.

Recovery, using log files, requires marking the start and end of every transaction; if the log file records show a transaction start without the corresponding transaction end, recovery processes can recognize the transaction as incomplete. Such markers correspond to what are called *checkpoints* in the program design.

52.4.4.2 Roll-Forward Recovery. Another approach to recovery is to start from a known good state and redo all transactions that are known to have been accomplished correctly. This roll-forward recovery requires a database backup and a transaction log file to be synchronized, so that the first record in the transaction log file represents the first transaction immediately following a database backup. With these data in hand, it is possible to reconstruct all the modifications up to, but not including, the last complete transaction. All incomplete transactions are discarded, although the recovery program typically prints out all the details available for transactions that were not completed.

52.5 BATCH FILES. Operating systems have their own facilities for writing and executing batch files, sometimes called *scripts*. Batch files can be complex enough that they look like programs, with logical tests, branching, and so on.

52.5.1 Backup File Creation. The protection of files that are updated by batch processing programs is automatic, because the file is completely recopied and the original is left in the form it was in before the program executed. Therefore, each processing cycle leaves its own backup. The name *batch processing* comes from the idea of sequentially processing transactions in a group. There are always two or more files involved. There is the master file that is being updated, and one or more transaction files that contain the information used to update the master. The process copies master file records that have no update activity into a new master file, updates the master records that do have activity and copies them into the new master file, and does not copy those master records that are flagged as deletions in the activity file. When the process completes, there are at least three files: the original master, the activity file(s), and the new master. Backups are the original master and the activity file(s). In the next processing cycle, the new master becomes input, together with that cycle's activity file(s). If a problem is encountered with the master file, the prior cycle's master file and all subsequent activity files can be used to generate an updated master. Keeping two or three generations of master and activity files is a common practice. The term *generation* is often applied to the batch processing cycles—generation 0, generation 1, generation 2—or as *grandfather*, *father*, and *son*. In recent decades, these latter terms have been replaced by *grandparent*, *parent*, and *child*.

52.5.2 Audit Controls. Another security measure applied when working with batch files involves using control totals to assure that the batch process has executed with accuracy. Specific controls may include counts of records. There also may be control totals of the values in specific fields of the input records to compare with totals in the output, after processing additions and subtractions. For example, in a payroll system, information from time cards is submitted to a computer for processing. Along with the time cards is a transmittal that contains a count of the time cards, and totals of all the straight time and overtime hours compiled by the department that originated the time cards. A procedure is run in which the control totals on the transmittal are checked against the totals actually input from the time cards, to assure correctness before initiating the payroll process. If discrepancies are encountered, they need to be investigated and corrected before taking the next step in processing the payroll.

52.6 DATA INTEGRITY AND VALIDATION. The following sections present a brief overview of validation; however, for more extensive coverage of this topic, see Chapter 47 in this *Handbook* on operations security and production controls.

52.6.1 Validation Controls. Whether an application updates its files by means of batch or online processing, or in some combination of both, validation of input is paramount. The term *GIGO* (garbage in, garbage out) is not heard as often as it used to be, yet it is as true as ever. Most contemporary systems use online files and databases, with information entered interactively, allowing for validation at the source. There are several specific validation techniques that have been applied to reduce the amount of incorrect information input into files and databases.

52.6.1.1 Methods for Identifying Input Errors and Unauthorized Modifications

Data input errors are common.

- Fours and nines, ones and sevens, and sixes and zeros can easily be mistaken for one another.
- Operators can leave out a digit or insert an extraneous one.
- Transposition is another type of mistake that is made from time to time. For example, an order entry operator may key in 3286 instead of 3826 (a transposition error), or perhaps 7790 instead of 7796 (a transcription error).
- Such errors would be reported when the *check digit* calculation produces a number that differs from the input check digit.
- A check digit is an extra digit that is added to a generated variable as a suffix.
- For example, a five-digit stock number could have an extra digit; the sixth digit or *check digit*, as it is called, is calculated by applying an algorithm (performing some arithmetic) on the first five digits, resulting in a single-digit result.
- To minimize input errors, check digits can be used on any variables that are program generated; for example, part numbers, employee numbers, and codes of various types.

A single check digit can sometimes conceal the existence of double or multiple input errors in an input string because of the relatively simple arithmetic schemes. More complex versions of the check digit generate a sequence of digits called a *checksum*, which provides greater power to help identify input errors or fraud; credit card numbers typically include a four-digit checksum at the end of the number. For example, a typical VISA credit card number has the format 1111 2222 3333 4444; the fourth block (4444 in our example) is calculated using a secret algorithm as a function of the other three blocks. Computer criminals have often tried (and sometimes succeeded) to reverse engineer the checksum algorithm to create VISA numbers that actually belong to unknown victims. Credit card security organizations have learned to apply complex *heuristic methods* to identify unlikely transactions; e.g., if the legitimate owner has just filled up her gas tank at an often-used filling station in Vermont, it is unlikely that a transaction five minutes later for 1,000 wool blankets at a manufacturing firm in Italy is legitimate, regardless of the accurate-looking VISA number.

An extension of the checksum, the *hash total*, is a common tool for identifying logical or physical data corruption. A hash total is simply the meaningless sum of numerical values, such as part numbers, rather than the meaningful sum of quantities. Recalculation of a hash total can usually indicate if the data have been corrupted since the last time the hash total was computed.

A digital signature is a cryptographically generated value based on an encryption key. Applying the same digital signature process to the data should generate an identical signature. In addition, as described in Chapters 7 and 37 in this *Handbook*, using the public key cryptosystem allows one to verify the authenticity as well as the integrity of signed data.

52.6.1.2 Range Checks. Range checks offer another way of validating information at its source. In situations where quantities and currency values are being entered, lower and upper limits can be established. For example, a range check could

be used on quantity ordered to ensure that the value entered is within specific bounds, such as no less than 10 or no more than 60 of a certain item. Although not completely error proof, such limits can at least eliminate obvious errors and reduce those not so blatant. Range checking with values of 10 and 40 will eliminate the extra 0 error, where an operator enters 100 or 600, or 1 or 4.

52.6.1.3 Validity Checks Using Tables. Using tables of values or valid codes is one of the best ways to ensure that only correct information is entered into files and databases. In contemporary files and databases, data type is specified as a property of the field during the design phase, and these properties can be used to make sure that text information is not entered into numeric fields and vice versa. Properties are also used to filter out invalid data configurations. Although impossible to filter out misspelled last names and other open-ended types of information, it is possible to filter out invalid state codes and other violations of standards. For example, a table containing the valid state codes in a country can be used to assure that none but those codes can be entered in the state field. Other such entities (e.g., department codes, product codes or types, and price classes) are primary candidates for table validation.

With tables, it is also possible to make combination tables. For example, if a certain product code fell within a specific range, then only four price classes might be allowable. To be concrete, suppose that when a product code falls within the 600 to 699 range, then the price class could only be R, PR, OC, and CD. Such tables are used to validate information as it is initially entered into files and databases. Entry of invalid information, whether intentional or inadvertent, is inhibited.

Tables of forbidden combinations are particularly important in real-time control systems for processes that can go seriously wrong as a result of bad inputs. For example, in December 1992, an explosion at a chemical plant in the Netherlands was traced to a data input error that could have been prevented by an edit check that verified a table of forbidden chemical mixtures. The following is an excerpt from a report published in RISKS 14.22 (January 4, 1993) by Meine van der Meulen of The Netherlands Organization for Applied Scientific Research, Department of Industrial Safety:

> The accident started with a typing error in a prescription by a Laboratory worker. Instead of tank 632 he typed tank 634. In tank 632 there was stored resin feed classic (UN-1268) and normally used in the batch process. In tank 634 DCDP (dicyclopentadiene) was stored. The operator, who had to check if the tank contents was [*sic*] equal with the prescription, filled the reactor with the wrong chemicals.
>
> … a heavy explosion occurred which caused the death of 3 firemen of the works fire brigade and injured 11 workers [including] 4 firemen of the works fire brigade. The damage was estimated at several [10s] of millions NL guilders.

Tables should generally be used only for information that is relatively static, because of the need for maintenance. Frequent table maintenance can lead to errors that have a cascading effect. If a table is incorrect, then all the information supposedly validated by that table could be incorrect throughout the files and databases.

52.6.1.4 Diagnostic Utilities. Production programs need to operate on verifiably correct data. Every production program should include diagnostic utilities to scan databases for logically impossible values or deviations from reality. For example, a diagnostic routine might check to be sure that every header total matches the total computed from the linked line item values. Because older data may have been accepted

under older range-validation rules, all records should be checked to ensure that they conform to the current range constraints. Similarly, any other logical constraints in an application system should be explicitly addressed in the diagnostic program. The output from the program should give all the details appropriate for identifying and correcting the incorrect data. For example, instead of simply indicating "BAD VALUE IN PRICE FIELD, RECORD 1234," the output should explain something like, "PRICE FOR ITEM 234 = $45.67 WHICH EXCEEDS CONFIGURED LIMIT OF $40.00 SHOWN IN INVENTORY MASTER RECORD 78999."

52.7 CONCLUDING REMARKS. In another application of the 80/20 rule, 20 percent of the design and development work is expended in handling 80 percent of the data (the normal and the correct) going through a system, whereas 80 percent of the design and development work is expended for handling 20 percent of the data (the errors, the exceptions, and the unusual situations) going through the system. It is safe to say that more effort should be expended on validation, and on backup and recovery procedures, than is spent on those processes that handle validated data and the routine tasks of application processing.

Good application controls involve data validation at all tiers of an architecture as well as careful consideration of the entitlements granted to users from an application, systems, and database perspective. It is not enough to focus only on user input and trust input from other systems. In today's complex systems there are too many areas where data can be modified either accidentally or maliciously to trust any input without validation.

In addition, organizations need to leverage all of the features of database systems to help not only maintain integrity, but understand data backup and recovery mechanisms. Adequate backup and recovery processes and procedures have to be put in place to handle system failures. These processes and procedures range from disaster recovery to single system or application recovery. To ensure that sites, enterprises, or applications can be brought back after failures requires that processes and procedures be put in place during the design and development process and kept up to date as the situation requires.

52.8 FURTHER READING

Basta, A., and M. Zgola. *Database Security*. Cengage Learning, 2011.

Chadli, M., and H. Coppier. *Command-control for Real-time Systems*. Wiley-ISTE, 2013.

Jones, C. *Software Engineering Best Practices: Lessons from Successful Projects in the Top Companies*. McGraw-Hill Osborne Media, 2009.

Kroenke, D. M., and D. Auer. *Database Processing*, 12th ed. Prentice-Hall, 2011.

Lublinksy, B., K. T. Smith, and A. Yakubovich. *Professional Hadoop Solutions*. Wrox, 2013.

Ottenheimer, D. *Big Data Security*. Wiley, 2014.

Post, G. V. *Database Management Systems: Designing and Building Business Applications*, 5th ed. Self-published, 2010. http://jerrypost.com/Books/DBBookSummary.html

Radvanovsky, R., and J. Brodsky, eds. *Handbook of SCADA/Control Systems Security*. CRC Press, 2013.

Romney, M. B., and P. J. Steinbart. *Accounting Information Systems*, 12th ed. Prentice-Hall, 2011.

Shelley, G. B., and H. J. Rosenblatt. *Systems Analysis and Design*, 10th ed. Cengage Learning, 2013.

Sommerville, I. *Software Engineering*, 9th ed. Addison Wesley, 2010.

Stair, R. M., and G. W. Reynolds. *Principles of Information Systems*, 11th ed. Cengage Learning, 2013.

Van Wyk, K. R., and M. G. Graf. *Enterprise Software Security: A Confluence of Disciplines*. Addison-Wesley Professional, 2014.

CHAPTER 53

MONITORING AND CONTROL SYSTEMS

Caleb S. Coggins and Diane E. Levine

53.1 INTRODUCTION **53·2**
- 53.1.1 Prevention, Detection, and Response 53·3
- 53.1.2 Controlling versus Monitoring 53·3
- 53.1.3 Control Loop 53·5
- 53.1.4 Defining the Scope and System Requirements 53·5

53.2 SYSTEM MODELS **53·6**
- 53.2.1 Internal, One-to-One, One-to-Many, and Distributed 53·6
- 53.2.2 Automation and the Human–Machine Interface 53·7
- 53.2.3 Snapshots versus Real Time 53·8
- 53.2.4 Memory Dumps 53·9

53.3 LOG MANAGEMENT **53·11**
- 53.3.1 Log Generation 53·11
- 53.3.2 Types of Log File Records 53·11
- 53.3.3 Automation and Resource Allocation 53·15
- 53.3.4 Log Record Security 53·15

53.4 GENERAL LOGGING CLASSES **53·17**
- 53.4.1 Overview 53·17
- 53.4.2 Process Flow and Job Scheduling 53·17
- 53.4.3 Network Connectivity 53·17
- 53.4.4 Environmental Concerns 53·18
- 53.4.5 System State 53·18
- 53.4.6 System Components 53·19
- 53.4.7 Process Activities 53·19
- 53.4.8 File System 53·20
- 53.4.9 Access Controls 53·20

53.5 DATA AGGREGATION AND REDUCTION **53·21**
- 53.5.1 Centralized Data Stores 53·21
- 53.5.2 Filtered Queries 53·21
- 53.5.3 Analyzing Log Records 53·21

53.6 NOTIFICATIONS AND REPORTING **53·22**
- 53.6.1 Alerts 53·22
- 53.6.2 Trend Analysis and Reporting 53·23
- 53.6.3 Dashboards 53·24

53.7 CHANGE AND SECURITY IMPLICATIONS **53·24**
- 53.7.1 Change Management 53·24
- 53.7.2 Configuration Protection 53·25
- 53.7.3 Performance Considerations 53·25

53.8 MONITORING AND CONTROL CHALLENGES **53·25**
- 53.8.1 Overview 53·25
- 53.8.2 Industrial Control Systems 53·26

53.8.3 Mobile Computing 53·27
53.8.4 Virtualization 53·27
53.8.5 Cloud Computing 53·28

53.9 SUMMARY 53·29
53.9.1 M&C Checklist 53·29

53.10 FURTHER READING 53·30

53.11 NOTES 53·31

53.1 INTRODUCTION. Monitoring and control (M&C) systems provide critical services to computing operations through prevention, detection, and response. When aligned with security policy and appropriately implemented, these systems enable organizations to detect security incidents, mitigate risk, and analyze trends. M&C systems may prevent persistent unauthorized login attempts from a remote location, detect performance issues on a storage network, or alert an incident responder to the attempted theft of intellectual property. Automation is a key factor for proactive, real-time, and scalable M&C system implementations. Although it is possible to review individual log files within a limited scope, manual log review remains impractical and inefficient for modern enterprises processing millions of events on a routine basis. Although the scope of M&C systems may vary in size and complexity, the central components remain the same: data collection, data reduction, and system response.

Data collection includes log file generation and storage. Computer systems use *log files* to store records of events. Logged information includes system time, relevant user accounts, system identifiers, and the actions performed. Programs write or append data to log files that are often stored in local directory structures and forwarded to central log collectors for aggregation.

Data reduction is the process of using predefined or on-demand rules and logic to generate relevant, filtered, and actionable views of aggregate log data. Exception reports such as failed or unauthorized login attempts filter out acceptable events, enabling analysts and automated alert processes to focus on anomalies for additional follow-up. Visual *dashboards,* including performance metrics and risk-ranked events, provide information managers and production operators with an enterprise view of current production activities that may impact service delivery objectives.

System response encompasses the actions resulting from the data reduction phase. Actions include alert notification, corrective actions, and recovery procedures. Depending on several factors, the response may be automated or involve feedback through a human–machine interface (HMI).

The following self-assessment questionnaire is based on the areas covered in this chapter; it identifies elements to review for an existing or planned M&C system:

- What is the purpose and goal of the M&C system, from a business and technical service perspective?
- What are the security, compliance, business, and legal requirements?
- What systems, components, or processes need to be monitored or controlled?
- What model will be used to monitor or control the target?
- Where will the data be stored?
- How will the data records be retained and secured?
- What kind of data reduction and reporting is required?
- Who or what will respond to detected events and incidents?
- What self-defense mechanisms exist for the M&C system?

53.1.1 Prevention, Detection, and Response. M&C system capabilities generally include the ability to prevent, detect, and respond to a wide number of events. However, not all events and challenges need to be met with costly control systems or dedicated sensors. Some activities may be more suited for monitoring than control. You may also want to alert on activities that may not be alarming on an individual basis but may be significant in aggregate or at specific thresholds. Decisions to implement and configure monitoring and control systems should be founded on a risk-based approach.

Prevention refers to the system's ability to prohibit an undesirable action. Common tools used in network security include in-line intrusion prevention systems (IPSs) and unified threat management (UTM) appliances. Deploying these systems at choke points, between enclaves, and at the network edge can restrict the proliferation of malware across an enterprise network as well as block Command and Control (C&C or C2) channels for botnet clients. In the physical-security realm, single-person mantraps prevent individuals from tailgating into secure facilities.[1]

Detection refers to the ability to identify attacks and other security concerns through sensory input. Systems rely on automated detection mechanisms and human observation. M&C systems must have reliable detection mechanisms in order to manage security events and provide meaningful metrics on the state of the information infrastructure. Proactive detection may involve the authorized use of a vulnerability scanner by security controls testers. Reactive detection may include the identification of source IP addresses and VPN users who participated in a distributed denial of service (DDoS) attack on an Internet-facing network.[2]

Response refers to the output of M&C systems. A monitoring system may detect remote attempts to enable FTP and Telnet on production servers, log the activity, and send an administrative alert to an operator or centralized correlation engine. A control system, however, would correct the configuration files and service states or prohibit the changes from occurring in the first place. Additional human response may occur after the administrative alert or dashboard refresh, such as escalation to an incident-response team for resolution and postmortem analysis.[3]

Failure to detect and report security breaches within mandated time frames may have legal ramifications, as many states require breach notification.[4] Implementing *defense in depth* (layered defenses) increases the likelihood of threat detection, incident prevention, and response. For example, AC power loss at a secure facility and a nonresponsive battery or generator backup may render a closed-circuit television (CCTV) system incapable of recording a physical attack. An additional layer, such as physical security guards, can provide a degraded level of monitoring while technicians restore power and monitoring systems to normal functionality.[5]

53.1.2 Controlling versus Monitoring. *Monitoring* refers to the observation, logging, and reporting of systems, components, and events. Data collected from monitoring systems may identify intruders, locate system vulnerabilities, confirm operational status, and provide confirmation that controls are online and in place. In a network environment, a vulnerability scanner can perform authenticated scans on specific devices to identify missing patches and configuration issues without necessarily modifying or controlling the targets. M&C systems operate in two modes, *continuous* and *batch*.

1. **Continuous mode** provides ongoing measurement and compliance assurance. Systems that operate in real-time, continuous mode present a more accurate, up-to-the-second view of an organization's operational environment.

2. **Batch mode** takes place at specific points in time. An annual audit often utilizes batch mode analysis, to determine the effectiveness of system controls and processes. The results provide limited assurance for a specific point in time or historical range based on available data.

Control refers to what we can manipulate and audit. Rather than merely observing and recording information about an unauthorized change in the operational environment, controls can restrict such activities based on predefined configurations and processes. Controls may be *prohibitive* or *corrective* in action.

- *Prohibitive* controls prevent the change from occurring.
- *Corrective* controls return or restore a control target to its approved state. A host intrusion prevention system (HIPS) may prevent modification to critical system files, while logging and alerting on failed attempts to do so.

Control systems ensure that the operational environment reflects stated policies and standards. The Control Objectives for Information and Related Technology (COBIT) section AI1.9 and ISO 17799 section 9.7 identify the need to monitor system access and use (see Chapters 44 and 54 in this *Handbook* for discussion of these standards).

Organizations that still depend on passwords for authentication should not only establish policies and draft procedures to meet those requirements but also need to implement technical controls to enforce compliance. If the standard includes a 15-character minimum password length, with specific complexity requirements and a mandated limit on the period of validity, technical password controls should prevent users and administrators from defining nonexpiring, short, trivial passwords such as *abcde*.[6]

Auditors inspect the operational environment to provide objective assurance of an organization's control environment. Typically, auditors evaluate a sample of control sets at scheduled intervals. These point-in-time, manual efforts provide insight into the infrastructure's degree of compliance. However, they do not provide the same level of efficiency as integrating automation into a framework that supports continuous information assurance.

The environment directly affects system security. Temperature, humidity, air quality, and other variables may introduce unscheduled system failures in a work environment. Such variables which may be outside our direct control also require monitoring. Computer system–management applications monitor internal components as well as environmental and equipment temperature and power-supply voltage levels. Most organizations may not be able to control the temperature of an overheating processor or a lack of electricity provided by the utility company, but sensory thresholds in a monitoring application may trigger automated scripts or notify local personnel to initiate a system shutdown. In these situations, the monitoring thresholds and alerts provide systems and operators the opportunity to gracefully stop running processes or revert processing to an alternate system or location.

Organizations decide to implement systems with monitoring or control capabilities for a number of reasons including perceived risk, financial costs, logistical difficulties, and business value. For organizations with an immature security program and limited visibility into the operational environment, monitoring can also provide a means of transition to future controls. For example, monitoring can provide baseline data and allow reports on an organization's resource utilization (e.g., disk space, Internet usage) and provide a basis for effective planning ("We will run out of disk space between

August and November of next year") as well as indicating areas for better awareness, training, education, and compliance efforts ("Four of our employees are collectively responsible for 60 percent of our Internet bandwidth utilization—and half of their activity involves watching pornographic videos during working hours!").

53.1.3 Control Loop. For most environments, humans remain in the control loop. The *control loop* consists of the *controller*, the *target system*, a two-way *communication path*, and the *transmitted data*.

The control loop is considered a *closed* loop when the process is automated and does not require human interaction; an *open* loop involves human interaction.

The requirements of the system determine the loop type. When gas-pipeline sensors detect a sudden decrease in pressure, an automated control system may need to perform an immediate lockdown to isolate the leak and send alert notifications. Leaving the loop open and waiting for human response may result in undesirable product loss and environmental and legal challenges. For operating-system patch deployment, an open loop would allow a system to interrogate hosts, determine patch levels, and report status for further review and action. A supervisor could then approve specific patches to particular systems after sufficient testing on noncritical, nonproduction systems and coordination with system stakeholders. Although fully automated software patching may seem like a viable process, it introduces additional risk (through system and process disruptions) and compatibility concerns. For industrial automation and controls, specific updates may not be certified for use and require follow-up with the vendor or verification on a support Website.

53.1.4 Defining the Scope and System Requirements. For M&C systems to operate effectively, management must clearly define the *scope* and the *system requirements*.

The *scope* specifically refers to the relevant processes and the people who interact with the computing environment. For example, to monitor the research and development (R&D) department for network communications regarding intellectual property, the scope may include IT/IS support staff, shared network resources, and existing processes to communicate with remote sites and external departments.

The *system requirements* refer to the capabilities necessary to perform the desired functions of M&C. In working through the requirements, it is important to remember that prevention mechanisms and controls in real-world environments can eventually fail or become obsolete due to technological changes. The requirements for monitoring R&D may include the need for real-time packet captures, an algorithm to identify intellectual property, automated alerts, and a reasonably secured collection device. Installing a network test access port (TAP) on the network uplink and then writing the captured traffic to a data collector for future analysis may seem logical. However, if a site survey identifies the use of undocumented cellular modems, unencrypted wireless hot spots, and personal email accounts for sensitive communications, the network TAP will fail to address the original goal. In this case, the M&C requirements expand to include signal limiters (Faraday caging), mandatory device restrictions, network access control (NAC), enforced IPSec communication policies, and Web filters. Periodic security awareness training and a policy refresh may also be necessary to ensure the organization clearly communicates expectations on the use of personal email and other electronic communications and devices.

We also need to take into consideration the capabilities of the M&C system, including the unintended consequences of unnecessary features and capabilities on a

deployed system. For example, a PC-based system used to manage a physical security badge system could also be used by an operator to attach USB devices, play video games, and connect to the Internet for peer-to-peer (P2P) file sharing. These issues can be partially mitigated through system security hardening, agent-based tools, and network design. Distribution panels connected to badge readers perform more specific, restrictive processing functions than a PC or commodity server; for panels, the attention would be directed toward physical access to the devices and communication path.

Compliance requirements and an organization's security program influence the implementation of M&C systems. Regulations may require access controls to restrict information to authorized users and sufficient monitoring to detect unauthorized attempts to gain access to information resources. In response to HIPAA, an organization may need to integrate additional log review procedures as well as appropriate measures to control all systems and processes that touch or potentially transfer personally identifiable information (PII). A public company needs adequate controls, to ensure accurate financial reporting. A security policy may reference the need to comply with Sarbanes-Oxley Section 404. The scope of an M&C system would then need to include computing resources, information process flows, and personnel with access to the financial information system. M&C systems should reflect the expectations of an organization's security policy and should fit into a relevant framework, such as COBIT. Utilizing a control framework provides logical management of M&C systems. The framework enables management to further understand enterprise processes and determine the need for risk mitigation or acceptance. PO9.7 in COBIT identifies control systems as part of the safeguard selection portion of risk management.

53.2 SYSTEM MODELS

53.2.1 Internal, One-to-One, One-to-Many, and Distributed. M&C systems fit into four general models: internal, one-to-one, one-to-many, and distributed. The type of model used depends on the complexity of the system and the type of information required. These models may include autonomous systems as well as those that require human interaction.

1. An *internal system* is one of the simplest forms of M&C; it monitors or controls itself. A single server or network device may monitor processor utilization or other system components to evaluate system performance issues. From a control perspective, an internal system may prohibit specific actions, such as unauthorized attempts to modify system files or to log on locally to the system.
2. The *one-to-one model* enables one system to monitor or control another independent system. When scaled out across a large environment, this model can become cumbersome and result in unnecessary duplication of IT services and processing. For example, an enterprise job scheduler monitoring a single financial application server can identify when files are available for processing. If other systems within the same environment are operating independent job schedulers, operators will require additional time to monitor and troubleshoot the duplicate functions. However, the one-to-one model can be quite useful in specific, high-availability situations that require failover. A network firewall, or critical server, would continually monitor and mirror the configuration and session information of the primary system. Once the primary system faults, the monitoring system would assume production functions, to avoid significant disruptions to the business environment.

3. The *one-to-many model* is most common in enterprise environments. A central monitor and control system remotely manages a number of systems. From a central point, all targets may be interrogated and updated, without the need for manual, one-to-one changes. Centralized M&C can also reduce cost by simplifying M&C functions of multiple resources at various sites. This model also improves audit efficiency, as fewer sources of information must be reviewed to validate controls and changes within an environment.
4. The *distributed* model involves sensors and controls dispersed throughout the environment. The control elements may operate independently or remotely with input from a central control system. Nodes may serve as a local proxy or feed information directly to a central collector for secure log management, reporting, and data reduction. For organizations attempting to manage heterogeneous devices, a distributed model enables staff to implement the best M&C components for specific systems while still maintaining the ability to aggregate data into a centralized repository. This level of aggregation provides a wider view of change in the infrastructure.

53.2.2 Automation and the Human–Machine Interface. Automation enables M&C systems to collect sensory information and to initiate a response without the need for human interaction. Many information systems operate round the clock and generate large volumes of event logs. Delegating to operations staff the tasks of manually reviewing millions of log entries and correlating them into meaningful information is impractical. The purpose of M&C systems is to obtain meaningful intelligence from target systems and to report and correct anomalies. By automating data collection and reduction, managers and staff can review consolidated event summaries and respond to situations that have not already been automatically filtered or controlled.

The *human–machine interface* (HMI) is the point at which an operator communicates with a monitoring or control system. In industrial environments, HMI middleware transforms system activities into interactive screens, with enough capabilities for staff to perform specific job functions. The presentation of data is an important factor for HMI. Operators must have the ability to perform their job functions and to maintain adequate situational awareness, without access to unnecessary or distracting system elements. In common network operational environments, interactive Web interfaces serve as the HMI. An analyst may review stoplight coded alerts (red, yellow, green) in a Web interface and perform investigative steps to triage and resolve open issues.

In some situations, organizations leveraging automation may benefit from improved workforce safety and response time. Autonomous vehicles and devices continue to proliferate. The U.S. Department of the Navy utilizes unmanned undersea vehicles (UUVs)[7] which reflect the distributed model, performing autonomous actions, with information exchange through a sensor grid within the Global Information Grid (GIG).[8] These systems can monitor environmental variables, detect threats such as unexploded ordinance, respond to change, and communicate with system owners. Much like intrusion detection and prevention systems in a network environment, UUVs can perform functions that may be unsuitable for direct human interaction. Both can perform continuous monitoring and collect a large volume of information that could not be collected through manual efforts. Also, similar to intrusion prevention systems in a network environment, a UUV may detect a threat, alert nearby UUVs, and report status back via the GIG.

The expanding use of artificial intelligence (AI) and robotic resources provide additional opportunities for mobile and autonomous M&C. The U.S. Naval Research Laboratory's humanoid firefighting robot, SAFFiR, is another example of responding to environmental threats without exposing the traditional workforce to hazardous tasks while operating as part of a team. While SAFFiR may not replace a data center's gas suppression system, artificial intelligence within a data center has been seen in the fight against unsolicited commercial email and spam.

When deploying automated resources, implementing artificial intelligence, or remotely commanding unmanned vehicles, the integrity of the system and security of the communication paths are critical. Improper physical access, software buffer overflows, weak administrative credentials, and USB malware infections are common ways to compromise the integrity of an M&C system. In network environments, a network engineer may use out-of-band communication paths to connect to the HMI via a console port, in order to avoid common network attacks against a management interface. Using insecure communication methods to update or manage M&C systems across untrusted paths is another common way to expose M&Cs to unauthorized access and compromise. Migrating from FTP, Telnet, and password authentication to Secure Shell (SSHv2) with key pair authentication is one practical way to reduce some of this risk to a more acceptable level.

For network IPSs in a properly baselined network environment, wormlike traffic anomalies may be detected, routed to network tarpits, and cross-reported to additional sensors. In this situation, the automated response is significantly faster and potentially less disruptive than using an intrusion detection system (IDS; monitoring only) with alerts routing to an on-duty analyst for additional follow-up.

53.2.3 Snapshots versus Real Time. *Snapshots* provide a point-in-time view of a target system. Auditors use these for regularly scheduled audits. Scanning the environment on a monthly basis to confirm compliance can also unearth security trends within an environment. Trend reports identify security performance over time. An organization may notice a sharp decline in vulnerabilities after an external audit. Over time, without a stable vulnerability management program, the number of vulnerabilities will return to preaudit levels. Snapshots do not provide immediate or ongoing information on changes or corrective actions made to production systems. Additionally, snapshots will not verify that a control is working all of the time, only during the time that the snapshot was taken. Using monthly snapshots to review a system that frequently changes on a daily basis will not provide the same level of assurance as real-time M&C.

Real-time monitoring refers to persistent, ongoing observation of a target. *Real-time control* refers to a control system's ability actively to influence its target. Industrial environments depend on information gathered through real-time M&C. When liquefied natural gas (LNG) travels through an LNG terminal,[9] operators should not have to wait until the tanker is empty in order to confirm that any of the product successfully transferred to the storage tanks. An M&C system can continually monitor the volume of product leaving the vessel, compare that data to the volume entering the storage container, trigger alerts if a leak is detected, and initiate a shutoff to avoid product loss.

When allocating and distributing resources to meet demand, real-time monitoring provides overseers with an immediate understanding of the areas of need. In a network environment, this may include peak network usage during an accounting cycle or online shopping trends due to holidays. An Internet-facing shopping portal that uses cloud

computing for resource allocation can use real-time monitoring to transparently detect and provision or deprovision computing and storage resources based on customer demand. Internet sites facing distributed denial of service (DDoS) attacks may use real-time monitoring to identify the incident and real-time control to block or throttle undesirable traffic. A Web-caching service or content delivery network (CDN) may also be leveraged to minimize the impact. In industrial areas, real-time monitoring can detect where electric grid managers should route additional power resources to meet demand, how much to pull from reserves, as well as when to implement real-time controls to enforce rolling blackouts and ensure grid integrity.[10]

DS13.3 in COBIT 4.1 identifies the need for adequate event logging to analyze activities over time. However, without well-planned automated processes and data reduction, the volume of data generated from snapshots and real-time activities will overwhelm staff. Through centralized data collection and reduction, designated actors can evaluate the results of both real-time and snapshot reporting. This combination can identify general trends, ensure that M&C systems continually operate, and preserve evidence from an identified attack.

53.2.4 Memory Dumps. Memory dumps are representations of the data in memory at a specific point in time. Memory dumps are used most typically after a system failure, but they are also useful in forensic research when investigators want the maximum amount of information possible from the system. There are two approaches to obtaining copies of memory: online, using diagnostic utilities while the system is running, and offline, from magnetic storage media to which memory regions are copied.

53.2.4.1 Diagnostic Utilities. Diagnostic utilities are system software routines that that can be used for debugging purposes. Also known as *debug* utilities, these programs run at maximum privilege (*root* or *supervisor* level, or their equivalents) and allow the privileged user to see or modify any portion of memory. The utilities usually print or display the contents of memory regions in a variety of formats such as binary, octal (base 8), or hexadecimal (base 16), with conversion to ASCII for easier readability. The utilities generally can provide immediate access to memory structures such as terminal buffers that allow the analyst to see what specific users are typing or seeing on their screens, file buffers that contain data in transit to or from specific open files, spoolers (print buffers), and program-specific regions such as data stacks. Because debug utilities also allow modification of memory regions, they are to be used with the utmost circumspection; for security reasons, it is wise to formulate a policy that no debug utility with root access can be run without having two people present. In high-security operations, the output of debug utilities should be logged to paper files for proof that no unauthorized operations were carried out using these programs. Access to privileged debug programs should be tightly controlled, such as by strict access-control lists or even by encryption using restricted keys.

53.2.4.2 Output to Magnetic Media or Paper. One method of doing a memory dump and follow-up analysis is by copying the data from memory onto magnetic media such as tape, removable disks, or rewriteable DVDs. Although it was once practical to print the entire contents of memory to paper, the explosive growth of memory sizes makes such printing impractical in today's systems. In 1980, for example, a large multiuser minicomputer might have 1MB (megabyte) of RAM available, resulting in a manageable, half-inch-thick stack of paper. At the time of this

writing, it is more common to see business systems with 64-bit processors and over 4 GB (4,000 MB) of memory. For perspective, a relatively older PC with only 256 MB of memory would generate a single printout that could be several feet thick.

Copying memory to a static storage medium is, therefore, preferred for future analysis of memory dumps following system crashes, unknown malware infections, and incidents requiring forensic response.

53.2.4.3 Navigating the Dump Using Exploratory Utilities. On production systems using reliable operating systems, crashes are rare and generally explored thoroughly to identify the causes of the anomaly. Generally after creating a dump it is necessary to study it looking for the problem and then correcting it. For large-memory systems, exploratory utilities are used to speed the search for problems in the dump by allowing the analyst to find any named system table or other memory region. At a forensic level, hashing may also be used to filter for known-good and known-bad elements within the memory archive.

53.2.4.4 Understanding System Tables. System tables are directories of system data where each datum is identified by an assigned label, by its position in the table, or by pointers from other tables. An understanding of system tables is essential in analyzing system problems.

Regardless of the details of the operating system and the system-specific names of tables, some of the important system tables include:

- **Process control table.** Pointers to the process tables for each process that is running on the system or that was running when the copy of memory was obtained
- **Process tables.** Detailed information about each process, with pointers to all the tables for that particular process
- **Data stacks.** All the variables used by specific processes
- **Buffers.** Data in transit to or from files and devices, such as disks and terminals
- **Memory management tables.** Lists of available memory blocks
- **Inter-process communications tables.** For example, information about resources locking or any logical flags used by multiple processes

Working with someone who understands the detailed structure of the operating system tables can be critically important for security work in which investigators must determine exactly what happened during an intrusion or other unauthorized use of a system.

53.2.4.5 Security Considerations for Dump Data. Memory dumps must be secured while in use and destroyed when appropriate. The dump contains the totality of a system's information in memory, including such data as:

- Passwords that had just been typed into terminal buffers for use in changing logons or for accessing restricted subsystems and applications
- Encryption keys
- Confidential data obtained from restricted files and not authorized for visualization by operations staff (e.g., medical data from personnel files)

- Financial data that could be used in frauds
- National security information restricted to higher levels of clearance than the system administration

53.3 LOG MANAGEMENT

53.3.1 Log Generation. A *log file* contains a record of events. This record is the basic building block for M&C systems. These files serve as a digital audit trail for system activities. Operating systems, databases, and applications may contain built-in logging capabilities. However, logging is not always enabled or configured for use and aggregation. The system monitors activities and writes relevant information as a log file entry. Operating system log files track application issues, access attempts, and system-wide problems. Thresholds and file sizes can be adjusted to meet the requirements of the environment. It may be unnecessary to write a log entry every time a user successfully opens or closes a file on a network share hosted on a storage volume, but it may be necessary to log and report on privileged access attempts to a management console or remote connections via Secure Shell (SSH).

Transaction logs take logging one step further, by storing copies of the actual changes made to a system. They have a fixed size, automatically generating a new file for future transactions. Most common in databases, transaction logs can be used to rebuild a corrupt database. Using a known-good database backup and the transaction logs from the database backup to the present time, databases can be rolled forward to a reliable point in time prior to the failure. Transaction logs should be actively monitored, as they can quickly consume large amounts of storage. Normally, committed transaction logs are purged after a successful backup. Other log types, such as Website traffic, may automatically generate new log files once per day; these logs may not contain entire Website transactions, but they will identify the visitor based on predefined criteria. The key here is the flexibility in determining what will be monitored and what format will work best for log generation.

Data retention policies must clearly define how long data should be retained, according to external requirements and internally driven policy. Retention requirements could range from 1 to 6 years, while some data may need to be retained indefinitely on backup media due to a legal hold. Log file archiving and storage become problematic over time: system-level log files may reach maximum size and overwrite existing log data or trigger a system shutdown; centralized log management systems may slow to a crawl when querying a relational database for data to export to a flat file. Log collection and analysis systems must be scalable as event rates and volumes may increase or burst over time. Some organizations address these variances through incorporation of cloud-based logging and storage services, while others rely on more traditional, internal storage solutions.

53.3.2 Types of Log File Records. A *log file* (also known as an *audit trail*) traditionally contains records about transactions that have updated online files or otherwise represent events of interest (e.g., logging on to the system or initiating a process). Some database or application log files also may contain copies of records that were affected before or after the logged transaction. There are a number of different types of log file records and each contains a specific kind of information. Typically, all of the operating system log records are stored in files with a maximum configured size, and that are identified by monotonically increasing file numbers. When the file fills up, the

operating system opens another one with the next number in sequence as its identifier. Because log file records of different types have different sizes, on some operating systems the log file is defined with a variable record length.

53.3.2.1 System Boot. The system boot log record contains information on booting up, or starting, the machine. Since this is related to a specific activity and generally to a specific area of the hardware or media, information on the boot can prove helpful when the boot fails and analysis needs to take place.

53.3.2.2 System Shutdown. The system shutdown log record contains information on when the system was shut down and by whom. This information can be invaluable when attempting to analyze a problem or find a saboteur. On systems that include emergency shutdowns (e.g., by calling routines such as *suddendeath* when system parameters fall outside the range of allowable values), the shutdown record may contain specific information about the cause of the emergency shutdown.

From a security standpoint, trustworthy system boot and system shutdown records can prevent a malefactor from concealing a shutdown followed by unauthorized boot to a diagnostic subsystem that would allow file manipulations without log records to track the operations. The boot records would show an unexplained gap between shutdown and boot.

53.3.2.3 Process Initiation. A *process* begins when a specific program is loaded and run by a particular user at a particular time. Log records for process initiation show when the various processes were initiated and who initiated them. These files provide a method of tracking employee activity as well as monitoring events that occur. In addition, such records allow cost recovery using chargeback at different rates for different programs or for program launches at different times of day. More important, the record of which programs were executed by whom at which times can be invaluable in forensic research.

53.3.2.4 Process Termination. When reviewing the process termination log record, an administrator will be able to tell when each process completed or was terminated for some other reason. Some systems may provide more information, such as why an unscheduled or abrupt termination occurred, but not all process termination log records provide that information. Process termination records typically include valuable statistical information such as:

- Which process spawned or forked the process in question
- Identification of any processes spawned by the processes
- Number of milliseconds of CPU used
- Number of files opened and closed by the process
- Total number of input/output (I/O) operations completed by the process
- Total size of the memory partitions allocated to the process
- Maximum size of the data stack
- Number and maximum size of extra data segments in memory
- How many swaps to virtual memory were needed during the existence of the process
- Maximum priority assigned to the process for scheduling by the task manager

53.3.2.5 Session Initiation. A *session* consists of the communications between a particular user and a particular server during a particular time period. Whenever a user logs on to a system and initiates a session, a record of that event can be found in the session initiation log record. Frequent review of these records provides important information to alert administrators that an intruder is in the system. For instance, if an administrator knows that an authorized user is on vacation and the session initiation log file shows that a session was initiated for that particular user's ID, chances are significant that an unauthorized party used the system via a borrowed or stolen user ID. These records are particularly valuable during forensic work.

53.3.2.6 Session Termination. When a session terminates, for whatever reason, a copy of the time when it terminated is generally stored in the session termination log record. Much like the process termination record, the session termination record can include a great deal of aggregated information about the activities carried out during the session, such as total I/O, total number of processes launched and terminated, total number of files opened and closed, and so forth.

53.3.2.7 Invalid Logon Attempts. The invalid logon attempt file can prove invaluable in cases where logon attempts do not succeed. In some instances, the file can tell if the user attempted to log on with an incorrect password, if the user exceeded the allowed number of failed attempts, or if the user was attempting to log on at a seemingly unusual time. These log records can provide important information in cases where an administrator is attempting to track specific actions to a user or to ascertain if the logon failure was due to a simple error or to an attempted impersonation by an unwanted outsider. In hardwired networks, where every device has a unique identifier, the records usually include a specific identifier that allows administrators to track down the physical device used for the attempted logons.

53.3.2.8 File Open. The file open log record provides information on when each specific file was opened and by which process; in addition, the record generally records the mode in which the file was opened: for example, exclusive read and write, exclusive read, exclusive write with concurrent read, append only, or concurrent read and write.

53.3.2.9 File Close. The file close log record provides an administrator with information regarding when the file was closed, by which process, and by what means. The file usually captures information on whether the user specifically closed the file or whether some other type of interruption occurred. The records usually include details of total read and write operations, including how many physical blocks were transferred to accomplish the total number of logical I/O operations.

53.3.2.10 Invalid File Access Attempts. An important log record in the M&C effort, the invalid file access attempt shows the administrator when and to which files there were invalid file access attempts. The records generally show which process attempted the I/O and why it was refused by the file system (e.g., attempted write to a file opened for read-only access, violation of access control list, or violation of file-access barriers).

53.3.2.11 File I/O. Whenever information is placed into, read out of, or deleted from a file (input/output), the information regarding those changes is captured in the

file I/O log. As mentioned earlier, the log includes images of a record before and after it was accessed. The I/O log can be used in file recovery after system or application crashes. Coupled with transaction-initiation and termination records, such data can be used for automatic roll-back or roll-forward recovery systems.

The activities recorded here can prove especially helpful when trying to validate actions that were taken and to attribute them to specific individuals. Detailed logs are typical for databases, where a subsystem provides optional logging of all I/O. Application log records, designed and programmed by application developers, also typically allow administrators to enable such detailed logging.

53.3.2.12 System Console Activity. The system console activity file provides information on any actions that originate from or are viewed at the system console. Typically, the system console includes not only logon and logoff records but also special requests, such as printer form mounts, specific tape or cartridge mounts, comments sent to the console by batch jobs, and free-form communications from users. The console file records all such activity as well as every command from the system operator, and the system responses to those commands. These records provide an excellent tool for investigators tracking down the specific events in a computer incident.

53.3.2.13 Network Activity. Network activity files provide valuable information on activity taking place on the network. Depending on the sophistication and settings of the system an administrator is using, the information derived can be plentiful or scant. Specific devices may generate their own records on network activity; for example, routers, gateways, and firewalls may all keep their own log files. However, typically these are circular files in which records to be entered after the file is full are shifted to the start of the file, where they overwrite the oldest records. In forensic work, it is essential to capture such data before the information of interest is obliterated. Unfortunately, in many systems, the volume of network activity is so high that log files contain only the most recent minutes of traffic.

53.3.2.14 Resource Utilization. A review of the resource utilization log records will show all of the system's resources and the level of utilization for each. By monitoring this file, administrators frequently make important decisions regarding modifying system configuration or expanding the system.

53.3.2.15 Central Processing Unit. The CPU file shows the capacity and usage of the central processing unit for whichever system is being used and monitored. Based on this information, administrators can monitor when usage is heaviest and the CPU is most stressed, and can decide on utilization rules and requirements as well as on possible CPU upgrades. As in all log file analyses, any outliers (unusual values) and any unexpected change in usage can be investigated. Global CPU utilization records can be compared with the sum of CPU usage collected from process termination records; discrepancies may indicate stealth operation of unauthorized processes, such as malicious software.

53.3.2.16 Disk Space. The log records for disk space show the amount of disk space originally available on a system, the amount of disk space used (and generally what type of files it is being used for), and the amount of disk space that remains free and available. Comparison of total disk space utilization with the total space allocated

to all files can reveal problems such as *lost disk space* (i.e., space allocated to files that were never closed properly; such space is unusable by the file system because there are no pointers indicating that the disk sectors are actually supposed to be free). Such unallocated sectors or clusters also may be where malefactors hide data they have stored without authorization by bypassing the file system.

53.3.2.17 Memory Consumption. Important information on the amount of memory in a system and the amount actually being used can be obtained from the memory consumption log records. Such records typically also include information on virtual memory usage. Therefore, they can provide warning of *thrashing* conditions, where memory segments are being copied to disk and read back from disk too often. Details of memory consumption may be useful in tracking down unauthorized processes such as worms.

53.3.2.18 System Level versus Job Level. Different levels of monitoring, logging, and auditing may take place. Some of this may be preprogrammed, but in most environments it is possible to select the log records and related reports that are desired. For instance, a regular system level audit log may be produced at intervals, just to provide information that the system is up and running without any extraordinary problems. But an in-depth job level audit may take place more frequently in order to monitor the specific job and ensure that it is running correctly and providing the information needed.

53.3.3 Automation and Resource Allocation. Utilizing secure, well-organized log management systems to monitor and respond to events can lower the administrative effort required to manage an infrastructure. With the growing demands of compliance and costs associated with data retention, it is important to include log management within the scope of the enterprise data retention policy. E-discovery laws and postmortem attack analysis rely on efficient access to accurate log information. Logs require system resources, ranging from media storage to CPU cycles and memory utilization. Trade-offs must be considered when setting aside potential business-generating resources for the purpose of logging and reporting.

Manually reviewing millions of log files on various systems, in different time zones, and inconsistent formats is unrealistic for staff. For geographically dispersed sites, real-time log transfers may be delayed due to limited bandwidth over slow links. Not all logged information is of equal value. Some data records may tie directly to compliance issues, while others provide noncritical benefits to daily operations. Based on policy, management must determine what types of data must be aggregated and reduced in order to monitor and to ensure ongoing compliance. This is part of the scope definition process of M&C systems. Planners and implementers must not only scope current requirements but must ensure that the log management plan will scale to address perceived resource demands for operational, security, and compliance requirements.

53.3.4 Log Record Security. An organization must protect its log records from unauthorized access and modification. The four most common methods of log record security are: access control lists, checksums, encryption, and digital signatures. Combining these methods, based on the system requirements, can adequately protect most log data. Attackers with unfettered access to unencrypted email log files could download copies of the files and search them for sensitive information as well as

modify or delete message content directly from the target server. This type of attack threatens every element of the Parkerian hexad.[11] In addition to the common concerns of confidentiality breaches and loss of data integrity, unauthorized access to log files may provide an attacker with additional information on the network environment. Vulnerability scanners may run at scheduled intervals and log detected issues to a central database. Leveraging aggregated data, an attacker can more efficiently compromise a network environment and evade detection routines accustomed to high-footprint reconnaissance.

Operating system and application-level access control lists (ACLs) restrict access to files. Some log files may be written to write-once, ready-many (WORM) media for auditors or legal staff. Operations staff may have no legitimate need to access such data. However, they may need to ensure that the communication flow operates continually. Databases can implement even more granular controls on specific records and fields, to avoid unnecessary information disclosure. Physical access controls to the media while in storage or in transit must also be considered when planning an M&C system for log data. Properly implemented ACLs can limit the level of log data exposure to authorized personnel and system processes.

Organizations subject to PCI DSS are required to implement integrity controls on log data. Historically, 32-bit cyclic redundancy checks (CRCs) generated checksums to guard against file alteration. Today, stronger cryptographic checksums, such as those generated by the SHA-2 family, provide a more reasonable level of assurance.[12] Checksums do not require constant human interaction. For tamper protection, a monitoring system may baseline a system and store copies of files or their hash equivalents for periodic or real-time file comparison. Changes to logs or other critical files could then be centrally reported for manual investigation. A control system could automatically quarantine the unauthorized file and restore a copy of the original file.

Digital signatures use public key cryptography to provide data integrity and nonrepudiation. These build on standard cryptographic checksums because they not only confirm file integrity but also identify who or what entity is providing that assurance. For organizations that must submit reports to external recipients, such as government offices, digital signatures can decrease the paper trail and provide a more efficient, digital means of meeting regulatory requirements.

Encryption must be considered for log data in transit and at rest. Encrypting entire log files reduces the ability to view or alter the contents. It also increases the resource usage on the target system during the encryption process. Like checksums, the keys required for decryption should be stored in a secure location, to increase the difficulty in defeating the control. Utilizing whole-disk encryption and then transferring a log file over a network in plaintext is not a consistent method of log security. Establishing secure network links, such as implementing IPSec, provides end-to-end link security. A combination of network and system-level encryption may be necessary to secure log files and other sensitive information properly.

The chain of custody should also be considered when planning a log management system. Some log data may need to be used in court. Without a verifiable audit trail, the contents may be inadmissible. A documented chain of custody for media transfers between staff and sites can validate physical controls. Implementing cipher block chaining (CBC) can further protect against unauthorized changes to log data. The previous log file ciphertext would Exclusive-Or (XOR) with the current plaintext to create a new value for encryption operations. This concept can also be useful for record sets within log files and database applications, as it will identify record deletions and

other data modifications. See Chapter 7 in this *Handbook* for more information on cryptography.

53.4 GENERAL LOGGING CLASSES

53.4.1 Overview. M&C activities vary, based on the business scope, objectives, and target system functions. Targets can be grouped into general classes including: process flow, job scheduling, network connectivity, environmental measurement, system state assurance, system component status, process activities, configuration settings, file system information, and access control. Current regulations may include only a subset of an organization's information systems in the scope of compliance. Changes to the infrastructure and compliance requirements are inevitable. It is important to implement scalable and flexible M&C in order to accommodate future needs.

53.4.2 Process Flow and Job Scheduling. Some organizations use a combination of mainframe and distributed system resources. The mainframe group may use a particular batch job scheduler to process and monitor job status and output. Other distributed systems throughout the organization might use their own schedulers and logging capabilities. The complexity of this design, however, does not facilitate enterprise-wide monitoring and auditing. A centralized job scheduler may be necessary to monitor and control job flow. The centralized system will need to connect to disparate systems to collect job status and initiate commands. For management, an enterprise view of job scheduling can identify unnecessary complexity and redundancy, leading to future process improvement. When shifting to a central job scheduling system, it may seem desirable to shift the entire environment immediately or possibly to revert to a central mainframe. However, a gradual migration to centralized M&C would be a more ideal method of change.

53.4.3 Network Connectivity. *Network connectivity* refers to the devices, protocols, and communication media (wired or wireless) used in a computing environment. A network operations center will want to monitor the status of network links, health indicators of critical devices, and data streams traversing the network. In a layered environment, physical access to the network cabling may be secured in cable raceways and network closets or at distribution points. Devices will be up to date and configured according to best practices. Unnecessary protocols will be controlled by blocking or routing suspicious traffic to network black holes and honeynets for further analysis. Network intrusion prevention systems (IPSs) can identify a communications anomaly, block the communication exchange, log the event, and trigger an administrative alert. Prevention requires planning and sufficient controls to deny an attack from achieving its goal. A method to implement M&C for sensitive network nodes may involve moving nodes to isolated segments, deploying in-line network IPSs, and utilizing host-based security agents. Additional enterprise security assessment tools may also scan, report, generate tickets, propose remediation measures, and analyze security trends over time.

In addition to monitoring general network connectivity, some networks are specifically deployed for the purpose of M&C. Legacy industrial systems now integrate common Ethernet networking elements into mechanical, electromechanical, and electronic environments. Wireless M&C options are also available. The ZigBee[13] specification utilizes the IEEE 802.15.4 standard to provide ad hoc local networking connectivity.

ZigBee is primarily used for M&C. In a plant floor or outdoor facility, additional cabling for M&C systems may be infeasible. A potential solution may be to integrate inexpensive, robust ZigBee control points. For geographically dispersed sites, high-bandwidth, secure connection services such as satellite, leased-line, and virtual private (VPN) network connections bring sites together to manage M&C data centrally.

53.4.4 Environmental Concerns. Some environmental variables are easier to monitor than control. Whether we are using ZigBee to transmit sensory information or a software package to measure AC voltage irregularities, environmental measures provide additional situational awareness for information systems and staff. Humidity, temperature, smoke, and leaking water directly impact computing systems in a data center. Dedicated air-conditioning units and dehumidifiers can monitor and control humidity and temperature. Indeterminate sources of smoke or water leakage may be difficult to control with automated response, but monitoring and alert notification may trigger corrective procedures by staff in order to mitigate risk.

Environmental factors may also impact the quality of an organization's products. M&C systems can enhance an organization's quality control efforts, ensuring consistent product manufacturing within predefined tolerances.

As part of a business-continuity plan, it may become necessary for an organization to monitor and prepare for environmental threats such as floods, structural collapse, and power loss. A monitoring system, with early warning capability, could notify staff to report to an alternate work site or to prepare for contingency operations. In the event that an environmental threat disrupts control systems, staff must be capable of manual override. The EPA recommends familiarity with manual operations, in the event of a supervisory control and data acquisition (SCADA) system failure.[14] In addition to manual override, integrating redundant M&C systems with seamless failover can decrease the likelihood of operational disruptions.

53.4.5 System State. The *system state* refers to a collection of critical variables on the target system. A monitoring system may need to track running processes in memory, services, open port connections, system files, hardware status, and log data. In industrial control systems, the system state may reflect the overall picture of an electrical grid, or it may reduce to whether or not a set of valves are open or closed. When an antivirus application process fails to start, resulting in a subsequent malware infection, the monitoring system can trigger an alert so that staff or an automated control system can isolate the infection. One method of monitoring the system state is to run an *agent* on the host system. Software agents monitor a predefined list of items and report status on a regular basis to a central monitoring hub. When the agent fails to respond or "check in," the monitoring system can issue an alert to investigate the issue.

Host intrusion prevention systems (HIPS) protect the system state of nodes in a network environment. HIPS analyze activities based on signature and heuristic detection routines. They provide a reasonable means of Monitoring and Control (M&C), with the potential for centralized reporting and attack correlation. For postmortem analysis, memory dumps can be useful in analyzing problems with the system state. If a network monitoring system detected botnet-related activity originating from a system in human resources, capturing the live system state may be advantageous to an investigation. As part of the evidence collection process of forensics, controlling and preserving the system state requires specialized training and tools.[15]

53.4.6 System Components. Monitoring applications can track the usage and overall performance of system components, including CPU, memory, and storage media. Operating systems also have built-in tools to monitor components and store results in log files. Inventory management systems collect component-level information as well as software information as configured by operations staff. Six months after an equipment refresh, an audit may identify systems operating with only half of the originally purchased memory. The logged information will help to determine the scope of the issue.

Monitoring the performance of system components can also identify changes in business need or system abuse. With the deployment of a new Web portal, existing Web proxy servers may no longer be capable of handling the heavier load. Logged performance data would corroborate the need to expand the infrastructure or add additional capacity to existing systems. Performance set points trigger alerts when components depart from normal operating levels. The alerts may identify externally driven changes, such as voltage drops, or internal limits, such as maximum memory utilization. Heavy resource utilization on a historically idle PC running HMI may even point to misuse of the system for non-SCADA-related activities, such as internal network scanning, password cracking tools, or gaming.

System component monitors most often attempt to identify overutilized CPUs, heavy memory consumption, and full storage media. However, identifying routinely underutilized resources gives management the ability to reallocate resources to areas of need rather than purchase unnecessary additions to the infrastructure. Disk quotas control the allowable amount of usable media for users and applications. To protect a system from unintentional or malicious overutilization, quotas should be reviewed as part of a layered strategy. Virtual machines use storage, memory, and processor quotas to control the resources utilized on a host system. With large virtualization environments, an enterprise virtual manager system may be necessary to ensure adequate load balancing across all of the virtual system hosts as well as to determine if additional tuning is required to improve virtual machine performance.

53.4.7 Process Activities. A *process* refers to a running program in a computing device. Processes may initiate at the system or user level and may contain multiple threads. M&C processes can prevent malware infections and identify abnormal process activities. Antivirus systems typically address this area. The SQL Slammer worm targeted a vulnerable process over a network connection, injecting code into memory without writing to disk storage.[16] In addition to process monitoring, network traffic analysis and network IPSs could mitigate this type of activity. Using enterprise monitoring tools to aggregate critical process information on a number of systems will help to identify application issues, such as runaway processes. For performance tuning and resource evaluation, agents can also monitor the CPU time of running programs for high utilization.

Monitoring process activities can also play a role in software metering and reporting. Organizations may want to identify application usage for different departments, in order to avoid deploying unnecessary applications and paying licensing fees for *shelfware*. Agents can centrally report elapsed time for target processes, for data aggregation, and for management report generation. Monitoring the use of specific types of software utilities, such as those used primarily by administrative and security personnel (e.g., nmap, tcpdump, hping, Wireshark, and Microsoft SysInternals), could also indicate unauthorized network reconnaissance. Blocking the execution of known applications

by file name or hash value can limit the level of follow-up required to locate and kill remote processes. Monitoring for uncategorized or unapproved processes may further help in identifying the cause of performance degradation on remote systems.

53.4.8 File System. File system activities provide a wealth of auditing information. File systems store log data ranging from application errors to unauthorized authentication attempts. Just like IDS logging, file system logging requires performance tuning to aggregate the data and to avoid data loss. Logging every file access attempt and data modification on a system will fill up event logs with generally useless information. If system owners configure the log file to overwrite when full, actual critical events may be lost. The volume of log entries also impacts network resources. A constant flow of file system activity logs may negatively impact the performance of the log collector. Focusing on critical system files, sensitive data, and configuration settings will decrease the number of log entries and provide a more targeted set of results.

The configuration files for operating systems and applications require layers of protection to avoid unauthorized modification. Much like the layers discussed in Section 53.3.4, regarding log-record security, Access Control Lists (ACLs), encryption, digital signatures, and checksums can play a critical role in protecting objects within the file system. Attempts to bypass M&C systems may begin with configuration files and settings. End users attempting to bypass Internet Web filters may reconfigure their Web browsers to use a rogue proxy server. Attackers attempting to modify the configuration of a poorly secured Web server may leverage the security lapses to target additional systems connected to the internal network. Controlling changes to configuration settings via system policies and HIPS will mitigate these kinds of attacks.

53.4.9 Access Controls. Access control is a central part of security management and compliance efforts. Organizations often rely on more than one stand-alone system for identification and authentication, including applications, network operating systems, and building security. As a result, they face significant challenges when attempting to communicate, aggregate, and manage access control events. Adding and removing staff to a secure environment may involve badge identification, security tokens, and passwords. An M&C system targeting access control will need to ensure that unauthorized personnel do not have access to any physical or digital resources. Management may first need to develop a process flow with support from operations staff and business stakeholders.

Aggregating access logs and detailing failed login attempts and multiple user instances aid in the identification of system abuse. When users log on to network resources, they establish a *session* with the target system. Multiple sessions for the same user, at multiple sites, may indicate that a user account has been compromised or is being shared. Depending on the level of importance, this information may surface in an exception report, trigger an administrative alert, or result in an automated account lockout. Similarly, a series of invalid login attempts may reach a threshold and automatically disable the user account for a predefined period of time. This control minimizes the effort required to actively monitor system accounts while ensuring that attackers are not able to perform a brute-force attack on the credentials of an authorized user.

See Chapters 28 and 29 in this *Handbook* for further discussion of identification and authentication.

53.5 DATA AGGREGATION AND REDUCTION

53.5.1 Centralized Data Stores. Systems generate logs in a variety of formats that are not always easy to access. The first phase in data aggregation is ensuring that logging is working properly on the target system. This may require installation of additional log collection software, in order to push or pull data in a standard *syslog* format. The second phase, gradually connecting logging points to a centralized data store, builds on individual logging mechanisms to create an enterprise view of events. Centralized data storage provides a single point of analysis and audit, without the time-consuming process of manually interrogating individual systems throughout the environment.

Log data must be communicated in a secure manner. An agent-based system needs to establish a secure communication channel with the primary collector, to avoid distributing log details to an unauthorized collector. It may also be advantageous to prohibit the log aggregator from collecting unsolicited syslog messages. This security measure may be built into the log forwarding or aggregating mechanism, or it may utilize IPSec over controlled network links.

A major problem with large volume data stores is processing time. Log processing can potentially delay or otherwise inhibit incident response. Depending on the resources available, it may be more economically advantageous to store secured log data in compressed archives while using agent-based monitoring only to report critical security issues or operational problems to central storage. The time and effort to log, parse, and report on every data block received by the central collector may exceed internal security-policy requirements and the perceived cost benefit. Filtering out noncritical events and alerts would then be part of the baseline process for the central M&C system. Storing and analyzing entire log files may serve as a secondary (slower) monitoring mechanism, while the security-specific alerts are handled on an immediate basis.

53.5.2 Filtered Queries. The volume of log file records can be overwhelming and generally impossible for operations staff to review individually and on a regular basis. Centrally storing the data does not change the infeasibility of the task. It does, however, lend itself toward improved data reduction. Transforming raw data into a meaningful, reduced form can minimize the noise and enable staff and automated processes to target higher-priority activities. Running filtered queries against log records can provide more immediate results on potential attacks or system issues. A scheduled query may list all failed logon attempts. Additional filtering may drill down to a set of computing assets (e.g., servers processing credit card data) or user types such as root-level accounts. Another more specific query may only list network traffic anomalies between the R&D department and an array of Internet proxy servers.

53.5.3 Analyzing Log Records. Log records can be valuable to the system administrator, but only if they are put to use. Properly using a log file requires monitoring and reviewing its records and then taking actions based on the findings. Such analysis typically uses utilities provided by the supplier of the operating system or by third parties, such as commercial software companies or freeware distributors.

53.5.3.1 Volume Considerations. Someone must make a decision regarding how big log files should be. Such a decision is generally made by the system administrator and is based on what is considered useful and manageable. If money is involved—for instance, if additional storage capacity is needed—the chief technology

officer (CTO), chief information officer (CIO), chief operating officer (COO), system auditor, and a variety of other stakeholders might be called on to participate in the decision-making process. However, in the three and a half decades from 1980 to the time of this writing in June 2013, disk-space costs fell about 37 percent per year compounded annually—from approximately $853,000 per gigabyte (in 2013 dollars) to approximately $0.04 per gigabyte in constant dollars, so disk space is not much of an issue anymore.[17] More important is that log files, like any file, may be left in an inconsistent state if the system crashes; closing a log file and opening a new one in the series is a prophylactic measure that ensures that less information will be lost in disk I/O buffers if there is a system crash. In addition, it is wise to put log files on a separate disk controlled by a separate processor; that way, one can disable write-behind (delayed writes) and force completion of I/O for every record of the log file.

53.5.3.2 Archiving Log Files. Archiving is a very important part of making and using log files. Every organization needs to decide how long to keep its log files and where to keep them. Decisions of this nature sometimes are made based on space considerations but more frequently are made based on legal requirements within specific industries. Some companies may decide to keep materials for a year or two, while others may be mandated by law to retain archived log files for seven or eight years.

The decisions regarding the archiving of log files should never be made arbitrarily; in every instance a careful review and check of company, legal, and industry requirements should be conducted. Based on the findings of that review, a written policy and procedures should be produced. Every organization needs to tell its employees what these policies and procedures are, so they can be aware of what materials are being kept, where they can be found, and how they can be accessed.

53.5.3.3 Platform-Specific Programs for Log File Analysis. Many operating systems are in existence, and although there are certain similarities in how they function, there are also differences. Special training on each specific operating system may be necessary in order to do a correct analysis of the platform-specific log files it generates.

53.5.3.4 Exception Reports. The volume of log file records can be overwhelming. For example, a customer-facing banking portal could generate log entries for logons, logoffs, financial transactions, and other record changes. It may also log attempts to penetrate Web application defenses or detect hijacked banking sessions. Exception reports allow the analyst to focus on specific characteristics (e.g., unusual wire transfer activities to foreign countries). Such tools greatly simplify detection and analysis of anomalies and are particularly helpful in tracking down unauthorized and undesirable activities.

53.6 NOTIFICATIONS AND REPORTING

53.6.1 Alerts. Log management must include layers of automation to reduce the time required to review and respond to events in the environment. How an M&C system deals with alerts will directly impact the security of the target system. Overwhelming operators with too many alerts will result in slower problem-response time as well as the potential to ignore critical issues. The purpose of alerts is to improve responsiveness to major issues, such as a leaking pipeline or multisite network outage. Alerts may be

configured at specific set points so that a responder will be notified only when the attack threshold is reached. Elements that are controlled in the environment may have a higher threshold for alert notification. A sudden influx of malware from the Internet may be immediately detected and filtered at the network edge. Notifying operations staff of a large volume of malware may be useful. However, spamming security administrators with alerts on every blocked attack would be an inefficient use of alert notification resources.

Additionally, the means of alert notification play an important role in business continuity and disaster recovery. Out-of-band monitoring systems may detect a network equipment error that could not otherwise be detected by the failing device itself. For alert recipients, multiple notification means may be necessary. Email, pager, cell phone, and SMS messaging are the most common forms. Alert visualization is another valuable means for operators actively monitoring the environment. Similar to the dashboard concept, visualization can reduce data collected from sensors into a visual map, with graphical indicators to identify system issues. A combination of dashboard-style alert monitoring and targeted message delivery may provide the best coverage for issues of high importance.

In industrial environments, operators utilize HMI tools to maintain situational awareness with field equipment. HMI uses virtual buttons and other visual elements to simplify industrial functions. If alerts are frequently generated for noncritical activities, staff will clear the alert list and ignore them over time. Ignoring alerts for critical issues can directly impact production and lead to costly repairs. In this situation, a management dashboard could display performance statistics that may trigger a field investigation into the root cause of declines in production performance.

53.6.2 Trend Analysis and Reporting. Data aggregation, reduction, and correlation lay the foundation for trend analysis. Data related to compliance requirements can be analyzed to determine the pace of security improvements within the organization. Further analysis may identify the level of consistency of an internal security program as well as the true drivers of audit compliance. If system vulnerabilities and business risk continually peak in between scheduled external audits, management will need to investigate the root cause of this unstable, cyclic behavior.

Chargeback refers to the ability to charge for the use of IT services. Some organizations include these costs in the overhead of doing business rather than billing individuals and departments. Regardless of the financial preference, enabling chargeback monitoring in an environment provides an understanding of how information resources are used. If the marketing department utilization has grown from 40 to 70 percent of the Internet bandwidth, and the helpdesk use has expanded from 30 to 50 percent of the cloud computing resources, the trends may encourage management to realign resources to meet business demand or investigate resource abuse. Monitoring resource utilization and demand can also benefit development and testing cycles, so that production resources are properly allocated.

Chargeback systems, such as internal billing services and those used by external service bureaus, prove valuable in monitoring and tracing system activities through their logs. In some states and cities, legislation requires that these types of services keep accurate logs for audit purposes. The original records are used not only for tracking but also in court cases as evidence. In most situations, both internal and external services are the secondary "proof" of transactions, while the initial system logs provide the primary proof.

The log file records provide the basis for sophisticated billing algorithms that encourage rational use of resources. In addition, because unexpected changes in expenses can stimulate interest by accounting staff and managers, billing systems can sometimes aide in the identification of unauthorized usage, based on historical patterns.

Exception reports focus on specific data anomalies, such as employee expense reports that suddenly exceed US$9,000. Managers may want to employ continuous monitoring of the financial system to evaluate exception report trends. In addition to misuse and fraud, exception reports can identify system problems. Generating exception reports for systems that regularly operate below acceptable performance levels may point to a failing component. Management, at an industrial site that regularly fails to account for missing product, may uncover procedural issues that result in wasted material or defective equipment that must be repaired to resolve the ongoing product loss. Exception reports may also identify an isolated issue or a widespread problem. If change management recorded the replacement of hardware on five pipeline segments and they all started leaking at the same time, aggregated log data and exception reports would support the need for hardware and procedural reviews. Or an exception report, based on sensor data collected in November 2002 along the Denali Fault,[18] may correlate with an environmental issue (e.g., earthquake) that would possibly explain the cause of an equipment malfunction.

See Chapter 10 in this *Handbook* for basic information about statistical methods.

53.6.3 Dashboards. *Dashboards* consolidate information into a relevant and easily understood summary. Text, visual, and auditory elements may be employed, depending on management needs and system capability. Dashboards build on Web-based technology with a database backend to present a high-level overview of operational status at near-real-time refresh rates. Charts, graphs, stoplight indicators (red, yellow, green), system visualizations, and alert banners enhance the dashboard experience when used in moderation. Employees typically reference dashboards from a standard Web browser or through large video screens mounted in command centers and common areas.

Department and management-specific dashboards should also be considered when implementing specialty M&C systems. Depending on the sensitivity of the information, rotating the dashboards on common screens or within the Web portal itself may provide an opportunity to increase employee awareness of organization-wide and department-specific performance. Associates may learn more about the current level of business risk, of current operational performance (e.g., system uptime, and product generated, or sales volume), or even of the current state of cybersecurity.[19]

In addition to dashboards, organizations continue to use a number of system management consoles to monitor and control specific elements within an infrastructure. System management consoles often cater more to technical operations staff than to upper management. System management consoles and dashboards will likely continue to coexist, due to functional needs.

53.7 CHANGE AND SECURITY IMPLICATIONS

53.7.1 Change Management. M&C systems enhance an organization's ability to identify and manage change in the information infrastructure. Environment isolation (Development, Test, & Production), source code control, and separation of duties are fundamental prerequisites for effective change management. Knowing when authorized and unauthorized changes occur can improve problem resolution and assist in

incident response. In a data center environment, centralized system monitors track operational status. When running configurations are updated on core network equipment, a monitoring system could alert site staff in the event that the changes were not authorized. The cause may trace back to the remote exploitation of a known vulnerability or to an unintended change by a network engineer connecting to the production environment instead of a test lab. In either case, the information reported enables management to assess the situation and to implement corrective actions to reduce the likelihood of future attacks or procedural oversights.[20]

53.7.2 Configuration Protection. When new systems are deployed or legacy systems are brought up to specific compliance levels, configuration concerns surface. Although policy-based standards may correctly define the appropriate configurations to protect the information assets, failure to consistently apply configuration templates or checklists introduces unnecessary complexity to vulnerability and configuration management. To address this challenge, M&C systems can identify anomalies and bring systems back into compliance using predefined criteria. The Payment Card Industry Data Security Standard (PCI DSS) requires file integrity detection and alerting. However, an organization may also find it advantageous to deploy a system that will automatically detect, alert, and correct system integrity discrepancies while reporting on the noncompliant files or settings.

In industrial control systems (ICSs), programmable logic controllers (PLCs) store and run ladder logic programs to control hardware devices in the field. Programmers or automation specialists update these programs using specialized software for industrial networks. Modifying the program will directly affect the output of the actuators connected to the PLC. For example, incorrectly modifying a tank sensor's set point to a value greater than the tank's capacity could result in product loss, requiring manual override. For that reason, the program must be protected from malicious and accidental alteration. Using a password to protect the running program and layered network defenses such as a network IPS contribute to configuration protection in this scenario.

53.7.3 Performance Considerations. Like all systems, M&C systems require resources to operate. Depending on the operational model, these resources may come from the same host system, a network connection, or a remote device. Failing to gauge the performance impact of a monitoring system could leave the operational environment in a degraded state. Additionally, deploying new systems during heavy loads in a business cycle may disrupt normal operations. Installing an enterprise file integrity monitoring system and performing a baseline on production systems during peak hours could lead to overutilization of system resources and could halt critical production jobs. Changes in M&C require planning and coordination with relevant business units and technical staff, to ensure that new systems meet both the business needs and the technical requirements to operate efficiently.

53.8 MONITORING AND CONTROL CHALLENGES

53.8.1 Overview. Legacy systems do not always possess sufficient M&C for today's business and regulatory requirements. The art of retrofitting security into a legacy system requires overcoming physical, technical, and psychosocial barriers. Physical barriers range from wired ICS installations to wireless mobile communications. Technical issues include learning to detect and monitor unique communication protocols, understanding the production performance requirements, and managing

systems that historically operated in isolated locations with limited security oversight. Psychosocial challenges are often the most difficult. Organizations and business sectors unaccustomed to mature M&C systems typically deny the need for such systems. Rationalization attempts generally include denial that an attacker would ever target their systems, claims that the business is running well, and concerns that additional changes or complexity will likely disrupt production. Security-awareness training and changes in leadership can partially address these misconceptions while pointing to the lack of M&C data available to support any claims that operations is running as efficiently and securely as it should.

53.8.2 Industrial Control Systems. Industrial control systems (ICSs) operate in a number of industries, including oil, water, and power. ICSs are generally divided into distributed control systems (DCSs) and supervisory control and data acquisition (SCADA) systems. DCSs are more commonly associated with processes such as oil refining. SCADA systems involve more human interaction via HMI, and they collect data from a number of remote devices. Some systems contain elements of both DCS and SCADA. Security issues have increased, due to network connectivity and risks to critical infrastructure. These systems are expected to continue to integrate on a global scale, with improved, centralized data aggregation and reporting. Management will find value not only in aggregating their own organization's data but in collaborating in industry forums to implement more stable, interoperable, and secure industrial M&C systems. Like mainframes, SCADA systems no longer operate in an isolated environment. Sites often communicate via leased lines, satellite, and dial-up modems. Some industrial sites utilize ZigBee wireless equipment rather than installing standard cabling to monitor and control remote devices. Technicians dial in with a modem or through a virtual private network (VPN) tunnel over the Internet to manage industrial systems remotely. Other direct connections from the Internet (via Remote Desktop or VNC) may also exist but are not recommended.

M&C challenges for these systems include authentication, network security, configuration protection, and physical security. Historically, authentication has not been required to use a PC to communicate directly with a PLC over an SCADA network. Some vendors now attempt to build authentication into the HMI. However, the network protocols used in SCADA (e.g., MODBUS/TCP, DNP3) can still send unsolicited messages to the PLC to execute arbitrary commands. Configuration protection is typically associated with the running program on a PLC, which may have a weak or nonexistent password. From the operator standpoint, HMI usually runs on an unpatched PC with application-level security controls. These controls are typically bypassed by attacking the improperly secured operating system or inserting an infected USB drive. Laptops used by engineers to manage and update PLC programs cross legacy air gap boundaries and are often connected to untrusted external networks, used for Internet and email, and exposed to off-site physical attacks. Physical security remains a challenge, as equipment may not be located in a secured area, or operators may lack adequate security awareness training to implement proper security procedures. Fortunately, these M&C systems may leverage existing information security resources to mitigate risk in a layered architecture. NIST SPEC PUB 800-82 clearly describes the security landscape surrounding these systems. Network firewalls, intrusion prevention systems, and more robust physical security would all contribute to a more secure SCADA environment. Vulnerability detection tools now integrate SCADA protocols, with the goal of detecting vulnerabilities and mitigating risk. Continuing to develop more secure

communication protocols, and establishing configuration protection guidelines, would further improve the authentication weaknesses in SCADA deployments.

M&Cs in an ICS environment become more complex due to ownership and responsibilities. Traditional IT organizations often support networks up to a network port but often lack the expertise and knowledge specific to industrial control environments. Control engineers may implement specific network elements based on production demands, without necessarily aligning configurations with security and policy requirements.

53.8.3 Mobile Computing. As we continue to integrate mobile and distributed computing systems, the complexity and difficulty required to secure systems increases. We have more variables to address. Information flows even further from the traditional network perimeter. Organizations continue to support bring your own device (BYOD) policies that introduce untrusted and insecure network nodes into the environment; these personal devices storing company data should be addressed through very clearly defined policies and standards. With advances in de-perimeterization, M&C systems must ensure that data remains secure while in transit and at rest, as determined by organizational policy. Mobile technologies communicate over wireless links, invisible to traditional M&C systems. Some wireless devices require centralized system management but rarely connect to the corporate network. Mobile encryption and mandatory VPN connectivity help to control data storage and transmission. With centralized policies, these systems promote more consistent data communications throughout an organization.

Radio-frequency identification (RFID) monitoring systems wirelessly identify and collect information from the target system. NIST 800-98 describes the technological benefits and risks related to implementing this technology. M&C systems often collect sensitive information. The privacy issues inherent in RFID will require thorough risk assessment to avoid potential data loss. Tagging all of an organization's employees with RFID badges that contain PII subject to HIPAA regulations would be unwise, as the RFID tags could potentially be accessed from over one meter away without the holder's knowledge. From an inventory perspective, RFID tags provide a means to automate site audits of company assets or to collect asset information without manually writing it down. To avoid compromise, the device standards use integrity checks (CRC) and require a password to render them nonresponsive.

53.8.4 Virtualization. Organizations use a variety of computing equipment to meet their business needs. Some equipment may remain underutilized, while critical database servers require additional storage space and working memory. Virtualization provides the ability to slice and pool computing resources into more efficient, targeted, and often portable business resources. Virtual systems generally consist of physical hardware, a virtualization interface (VI), and a virtual machine (VM). The VM may be an application or an entire operating system. Since the VI oversees the VMs, we must implement appropriate M&C at the VI layer as well as within and around the VMs.

Users of a virtualized operating system may not be able to detect that it is virtualized. For example, a Web hosting company may wish to provide dedicated virtual servers to its customers without the additional expense of dedicated hardware. In this situation, the organization could deploy FreeBSD servers and implement a degree of virtualization using "jails"[21] to provide a high level of compartmentalized, system-level access to customers. An operator with root access to a jailed environment should not see the processes executed by other jails nor the underlying system hosting the jails. However, the host system owner may monitor or control elements of the host system as well as

those within jailed environments. Additional monitoring of the host system may also be necessary, to detect jailbreak attempts.

Another type of virtualization, *paravirtualization*, utilizes a *hypervisor* as the VI between the hardware and the VM. Rather than create identical jailed environments, a hypervisor can be used to support more than one type of guest VM, such as Ubuntu, Fedora, and PC-BSD. This configuration is possible because the hypervisor handles all communications (hypercalls) between the hardware and the VMs. Along the same lines as a FreeBSD host with jails, the hypervisor can monitor and control activities within and around the VMs with varying degrees of transparency to the VM user. The capability of stealth VM monitoring makes virtualization a useful monitoring tool for honeypot research. For a greater level of system assurance, organizations will need tools to detect and prevent attacks against the hypervisor.

Portability, an advantage of virtualization, poses a challenge for M&C systems. VMs can migrate from one hardware device to another frequently, while the integrity of an off-line VM is unknown. In a managed, virtualized infrastructure, performance degradation issues detected in the VI may trigger VMs to migrate from one physical server to a less utilized resource. An unstable VM may also be manually forced off line and rapidly redeployed using a backup virtual disk, without disrupting other VMs sharing the same physical hardware. Tracking these moving targets requires reliable access to the virtual pool of resources, documentation mapping virtual to physical resources, as well as a clearly defined policy for virtual system provisioning. For example, an organization with policies that require clear separations between confidential and public systems should prohibit the same hypervisors from managing both payroll applications and public kiosk VMs on the same hosts.

53.8.5 Cloud Computing. Cloud computing uses several deployment and service models to provision flexible computing and storage resources. However, M&C in the cloud introduces additional challenges, due to the deployment variations and additional parties involved in functional tasks. In public clouds, the physical environment and network infrastructure M&C responsibilities belong to the hosting provider. As a result, organizations have significantly less visibility into daily network operations and incident handling activities, and must rely on third-party attestations and standard reports (e.g., SAS 70, SSAE 16). Decreased M&C visibility places increased reliance on contractual agreements between the cloud provider and the cloud tenant. Guidance published by ENISA and the Cloud Security Alliance can assist in the cost and risk benefit analyses.

In addition to standard computing and storage resource needs, organizations may also leverage cloud resources for log management, disaster recovery, and development environment isolation.

Security responsibilities and M&C capabilities are greater for an organization utilizing the infrastructure as a service (Iaas) model than other service models. Virtual system *instances* can be configured with standard logging capabilities at the operating system and application level. However, some of the hardware elements such as memory may be shared with neighboring cloud tenants. Platform (PaaS) and software as a service (SaaS) models provide customers with less visibility outside the application environment.

Management of cloud resources may be handled through a standard Web interface and APIs. While many of the monitoring and control elements may still exist, the approach to accessing and storing the data may vary. For example, an administrator may create a snapshot of a running instance without impacting the production system. Additionally, the snapshot may be mounted in cloud storage for forensic follow-up.

Controlling access to keys, management interfaces, virtual machines, and storage volumes are important targets for M&C in the cloud. Additionally, data encryption should also be considered where possible.

53.9 SUMMARY. Monitoring and control systems provide technology stakeholders with additional assurance related to system operations. M&Cs enable visibility into production operations and may also detect and prevent undesirable events and operational inefficiency. Controlling every element of an information system is generally impractical from a cost and risk-based standpoint. However, in taking a risk-based approach, we can identify what should be monitored, controlled, or accepted as a residual risk. Regulations, internal policies, industry frameworks, business requirements, and operational needs all contribute to the decision process in selecting and implementing M&C systems and functions.

Defined scope and business requirements drive the selection of an appropriate M&C model. While human interaction continues to play a critical role in legacy systems, increased sophistication in automation can lead to improved reporting capability and continuous assurance. Unique challenges related to automation surface in environments with specific requirements including industrial controls.

Monitoring, control, and auditing interrogate information systems, identify problems, initiate corrective measures, and generate meaningful information on the state of various systems and processes. The systems may involve human interaction or function in an automated, closed-loop format. Systems generate log files based on system activities, store them in secure locations, and enable aggregators to collect and reduce the data into actionable information. Although log aggregation improves the analysis process, the resources required to store and manipulate the data can become problematic. Not all events are useful. In order to manage reasonable data volumes, only specific activities should be logged or collected by the central management system. Log file data provides a means to analyze events historically at any measured point in time—a necessary part of infrastructure management. Organizations need to understand the value of assets, align business with regulatory requirements, and find new ways to leverage M&C to improve operational performance.

53.9.1 M&C Checklist. M&C systems may be reviewed at any phase of the lifecycle (e.g., planning, implementing, maintaining, and decommissioning). Assessing the appropriateness of M&C systems and components involves multiple actors. Enterprise Architecture may sit down and define the requirements and identify a short list of potential solutions, while Risk Management, Audit, and Information Security provide feedback based on their areas of focus. In larger environments, Risk Management may also consult with Legal, Senior Management, and IT Governance board members to ensure that M&C investments adequately align with business objectives, maintain acceptable risk levels, and integrate with the evolving maturity and direction of IT services.

Working through an M&C checklist should result in the identification of key elements for a current or future M&C implementation. The information may serve as an input to broader functions including risk management, information security, and enterprise planning. At a minimum, an M&C checklist should include:

- Identification of vision and goals for M&Cs
- Definition of how M&C will align with business
- Mapping of security, compliance, business, and legal requirements
- Scope, architecture, M&C model, and deliverables

- Components to be monitored or controlled
- Storage methods (store/forward)
- Data storage resources and security levels
- Data reduction and reporting outputs
- Response methods
- Self-defending M&C items (e.g., log blackouts/variance as a detective measure)

53.10 FURTHER READING

Badger, L., D. Bernstein, R. Bohn, F. de Vaulx, M. Hogan, J. Mao, J. Mesina, K. Mills, A. Sokol, J. Tong, F. Whiteside, and D. Leaf. U.S. Government Cloud Computing Technology Roadmap Volume I Release 1.0 (Draft) High-Priority Requirements to Further USG Agency Cloud Computing Adoption. NIST Special Publication SP 500-293. www.nist.gov/itl/cloud/upload/SP_500_293_volumeI-2.pdf

Badger, L., R. Bohn, S. Chu, M. Hogan, F. Liu, W. Kaufmann, J. Mao, J. Messina, K. Mills, A. Sokol, J. Tong, F. Whiteside, and D. Leaf. US Government Cloud Computing Technology Roadmap, Release 1.0 (Draft), Volume II Useful Information for Cloud Adopters. NIST Special Publication SP 500-293. www.nist.gov/itl/cloud/upload/SP_500_293_volumeII.pdf

Catteddu, D., and G. Hogben. "Cloud Computing: Benefits, Risks and Recommendations for Information Security." European Network and Information Security Agency (ENISA). November 20, 2009. www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment

Cloud Security Alliance. "Security Guidance for Critical Areas of Cloud Computing Version 3.0." November 14, 2011. https://cloudsecurityalliance.org/research/security-guidance

ISACA. Control Objectives for Information and Related Technology (COBIT) v5: http://www.isaca.org/cobit

ISO/IEC 17799:200527002:2005: "Information Technology—Security Techniques—Code of Practice for Information Security Management." www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=50297; http://www.iso.org/iso/iso_catalogue/catalogue_ics/catalogue_detail_ics.htm?csnumber=39612 (URL inactive).

Karygiannis, T., B. Eydt, G. Barber, L. Bunn, and T. Phillips. Guidelines for Securing Radio Frequency Identification (RFID) Systems. NIST Special Publication 800-98. U.S. Department of Commerce, April 2007, http://csrc.nist.gov/publications/nistpubs/800-98/SP800-98_RFID-2007.pdf

McKinney, D. U.S. Naval Research Laboratory, "NRL Designs Robot for Shipboard Firefighting.," March 7, 2012. U.S. Naval Research Laboratory. www.nrl.navy.mil/media/news-releases/2012/nrl-designs-robot-for-shipboard-firefighting

NIST Cloud Computing: Information Technology Laboratory. www.nist.gov/itl/cloud/index.cfm

Payment Card Industry Data Security Standard (PCI DSS): https://www.pcisecuritystandards.org

Sarbanes-Oxley Act of 2002. 15 U.S.C. 7201. Public Law 107-204. 107th Congress; www.sec.gov/about/laws/soa2002.pdf

Stouffer, K., J. Falco, and K. Scarfone. Guide to Industrial Control Systems (ICS) Security, Rev. 1. NIST Special Publication SP 800-82. U.S. Department of

Commerce, June, 2011. http://csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final.pdf

Ylonen, T., and C. Lonvick. "The Secure Shell (SSH) Protocol Architecture". The Internet Engineering Task Force (IETF) Network Working Group RFC 4251. http://datatracker.ietf.org/doc/rfc4251 and http://tools.ietf.org/html/rfc4251

53.11 NOTES

1. For more information on intrusion detection systems and unified threat management appliances, see Chapter 27 in this *Handbook*.
2. For more information about DoS attacks, see Chapter 18 in this *Handbook*.
3. For more information on incident response, see Chapter 56 in this *Handbook*.
4. Legislation on Notice of Security Breaches, National Conference of State Legislatures, www.ncsl.org/programs/lis/cip/priv/breach.htm
5. For more information on physical and facilities security, see Chapters 22 and 23 in this *Handbook*.
6. See Chapter 28 in this *Handbook* for discussion of identification and authentication.
7. U.S. Department of Navy, "The Navy Unmanned Undersea Vehicle (UUV) Master Plan," November 9, 2004, www.navy.mil/navydata/technology/uuvmp.pdf
8. Department of Defense Directive No. 8100.1, "Global Information Grid (GIG) Overarching Policy," September 19, 2002, www.acq.osd.mil/ie/bei/pm/ref-library/dodd/d81001p.pdf
9. Federal Energy Regulatory Commission, "Existing LNG Terminals," www.ferc.gov/industries/gas/indus-act/lng/exist-term/everett.asp
10. Consortium for Electric Reliability Technology Solutions, "Demand-Response, Spinning-Reserve Demonstration—A New Tool to Reduce the Likelihood of Outages," news release, July 26, 2006, http://certs.lbl.gov/press/press-7-26-06.html
11. See Chapter 3, page 3 in this *Handbook*.
12. K. Scarfone, M. Souppaya, A. Cody, and A. Orebaugh. NIST Special Publication 800-115, Technical Guide to Information Security Testing and Assessment, U.S. Department of Commerce, September 2008, http://csrc.nist.gov/publications/nistpubs/800-115/SP800-115.pdf Keccak was announced as the winner of the five-year SHA-3 competition on October 3, 2012: http://csrc.nist.gov/groups/ST/hash/sha-3
13. ZigBee Alliance: www.zigbee.org
14. U.S. Environmental Protection Agency, "Suggested Pre-Hurricane Activities for Water and Wastewater Facilities," http://water.epa.gov/infrastructure/watersecurity/emergencyinfo/pre-hurricane.cfm
15. Helix by e-Fense: www.e-fense.com/h3-enterprise.php
16. CERT® Advisory CA-2003-04 MS-SQL Server Worm, January 27, 2003, www.cert.org/advisories/CA-2003-04.html
17. An HP7925 120MB disk drive cost $25,000 in 1980, about $100,000 in 2013 currency. The cost per GB in 1980 was ~$853,000 in 2008 dollars. In 2013, a Western Digital Caviar Green 3 TB SATA III 64 MB Cache Bare/OEM Desktop Hard Drive—WD30EZRX for sale at consumer stores cost ~$123, or about $0.04/GB. The ratio of the change in cost is ~4 × 10–7. Taking the 33rd root (2013–1980) of the change gives an annual cost change of ~64 percent per year (37 percent decline per year compounded annually).

18. U.S. Geological Survey, Fact Sheet 017-03, "The USGS Earthquake Hazards Program in NEHRP—Investing in a Safer Future," by J. R. Filson, J. McCarthy, W. L. Ellsworth, and M. L. Zoback. Last modified May 17, 2005. http://pubs.usgs.gov/fs/2003/fs017-03
19. Talisker Security Wizardry, The Computer Network Defence Internet Operational Picture, http://securitywizardry.com/radar.htm
20. For more information on software development and quality assurance, see Chapter 39 in this *Handbook*; for information on managing software patches, see Chapter 40.
21. For more information on jails, see Chapter 15 of the FreeBSD Documentation Project's *FreeBSD Handbook*, contributed by Matteo Riondato: www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/jails.html

CHAPTER 54

SECURITY AUDITS

Donald Glass, Richard O. Moore III, Chris Davis, John Mason, David Gursky, James Thomas, Wendy Carr, M. E. Kabay, and Diane Levine

54.1 INTRODUCTION 54·1

54.2 ADDRESSING MULTIPLE REGULATIONS FOR INFORMATION SECURITY 54·3

- 54.2.1 Gramm-Leach-Bliley Act 54·5
- 54.2.2 Federal Information Systems Management Act 54·5
- 54.2.3 Health Insurance Portability and Accountability Act of 1996 54·7
- 54.2.4 Payment Card Industry Data Security Standards 54·7
- 54.2.5 Publicly Available Security Publications 54·8
- 54.2.6 Risk Framework 54·10

54.3 AUDITING STANDARDS 54·10

- 54.3.1 Introduction to ISO 54·10
- 54.3.2 ISO/IEC 27000 Series 54·11
- 54.3.3 AICPA Security Audit Standards 54·14
- 54.3.4 ISAE 3402 Standards 54·19

54.4 SARBANES-OXLEY ACT 54·20

- 54.4.1 Introduction 54·20
- 54.4.2 Section 404 54·20
- 54.4.3 Achieving Compliance 54·21
- 54.4.4 Audit and Certification 54·25
- 54.4.5 Sarbanes-Oxley Conclusion 54·25

54.5 TECHNICAL FRAMEWORKS FOR IT AUDITS 54·26

- 54.5.1 Framework 1: People, Processes, Tools, and Measures 54·26
- 54.5.2 Framework 2: STRIDE 54·26
- 54.5.3 Framework 3: PDIO 54·27
- 54.5.4 General Best Practices 54·27
- 54.5.5 Technical Frameworks Conclusion 54·28

54.6 FURTHER READING 54·28

54.7 NOTES 54·29

54.1 INTRODUCTION. Traditional auditing focused on reviewing organizational financial records and controls to validate the accuracy and integrity of financial data. External auditors typically focused on material or macro-level issues, and internal auditors focused primarily on transaction-level controls, protecting assets, and

validating information adequacy and validity. However, changes in the regulatory environment over the past several years have caused a significant increase in the scope and depth of these responsibilities.

With the changed environment, the extent, focus, reliance, and quality of the audit work has increased significantly in quantity, depth, and reliance. Thus, the current business environment, whether global or domestic, increasingly requires routine, periodic security audits, and inspections that include compliance with:

- Legal requirements; e.g., U.S. Federal laws such as Sarbanes-Oxley (SOX), the Health Insurance Portability and Accountability Act (HIPAA), and the Gramm-Leach-Bliley Act (GLBA)
- Industry standards such as the Payment Card Industry Data Security Standard (PCI DSS) for organizations accepting credit card payments
- Auditing Standards, such as ISO 17799, ISO 27001, and ISO 27002, Statement of Auditing Standards (SAS) 70 audits, both Type I and Type II; and Statement on Standards for Attestation Engagements (SSAE) No. 16

These audits typically address networks, host systems, Internet-related connectivity, protection and detection devices, physical premises, and policy-related items on a recurring basis, since a static posture may unwittingly leave the organization vulnerable.

These reviews and audits may be regarded by an organization's management with attitudes varying between two extremes:

1. They are a cost of doing business, one required by a regulator or governmental law.
2. It's nice to have, but don't spend too much money.

Both extremes, particularly the latter, ignore the value to the organization from such reviews, since the expense may be the primary (or only) criterion evaluated by senior management.

Sarah Abelow summarized the value of information technology (IT) audits as follows:

1. **Reduce risk.** IT audits that are planned and executed based on best practices will identify and assess the risks in an organization's IT environment. … Once risks are assessed, there can be clear vision on what course to take—to reduce or mitigate the risks through controls, to transfer the risk through insurance or to simply accept the risk as part of the operating environment …
2. **Strengthen controls (and improve security).** Based on the assessed risks as discussed above, the next step is to identify and assess corresponding controls. If controls are assessed to be poorly designed or ineffective, corrective action can be taken.
3. **Comply with regulations.** Wide ranging regulations at the federal and state levels include specific requirements for information security. The IT auditor serves a critical function in ensuring that specific requirements are met, risks are assessed and controls implemented. … It is critical to an organization to have assurance that all requirements are met.
4. **Facilitate communication between business and technology management.** An audit opens channels of communication between an organization's business and technology management. Auditors spend time interviewing, observing, and testing what is happening in reality and in practice. An audit can provide valuable information in written reports and oral presentations. Senior management needs to know firsthand how their organization is functioning.

The reverse communication from senior management to the technology professionals is also critical. Senior management's objectives and expectations are communicated to the technology staff. Auditors can facilitate this communication through participation in technical meetings and through review of the current implementations of policies, standards, and guidelines. …

5. **Improve IT Governance.** The IT Governance Institute (ITGI) provides this definition:
   "IT Governance is the responsibility of executives and board of directors, and consists of the leadership, organizational structures, and processes that ensure that the enterprise's IT sustains and extends the organization's strategies and objectives."
   The leadership, organizational structures, and processes mentioned in the definition point to IT auditors as key players. IT auditing and overall IT management are focused on the value, risks, and controls around an organization's technology environment. IT auditors review the value, risks, and controls in all of the key components of technology—applications, information, infrastructure, and people.[1]

In addition, IT and security audits can help ensure that the organization is not unwittingly posing an IT-related risk for or threat to other organizations, whether related, foreign, competitor, peer, or unrelated, by becoming involved, even unwittingly, in national security breaches, industrial espionage or sabotage, computer trespass, privacy violations, spam, or harassment because of inadequate controls.

Recent developments in government regulation and oversight have reduced auditors' gray areas; that is, the number of ambiguities and exceptions to requirements for reporting and compliance have been reduced. Perhaps as a result of increasing public awareness of leaks of personally identifiable information (PII), several undergraduate and graduate programs have been developed in universities and colleges around the world to train students in the discipline of IT, security, and compliance auditing.

Finally, it is useful to remind readers of the difference between an *audit* and an *assessment:*

- An *audit* is a formal process, usually carried out by a certified audit professional, focusing on verifying compliance with stated internal policies, external formal standards, and legal requirements. An audit report usually provides detailed lists of exactly which formal requirements have been met or not met.
- An *assessment* is an informal process, usually carried out by an industry expert who may or may not be certified by a professional body, focusing on improving efficiency and effectiveness, often using informal industry best practices and the expertise of the evaluator. An assessment report usually provides detailed analysis of failings and concrete recommendations for specific ways to improve what has been assessed.[2]

This chapter focuses on security audits.

54.2 ADDRESSING MULTIPLE REGULATIONS FOR INFORMATION SECURITY.

Federal regulations and industry standards increasingly recognize the need to address controls related to information security. However, because of continual advances in technology and the diversity of equipment and structures available for use, regulations and standards have continued to be vague in stating specifically how security configurations should be implemented, tested, and monitored. Although there are many different clients from diverse industries, with various regulations, legislation, or industry specific standards, all of them have a fundamental source of their best practices.

During the 1980s, the Defense Computer Security Center began to create a series of standards related to securing computer systems for governmental and subcontracted entities. The series of documents relating to the configuration, auditing, and securing of systems is commonly referred to as the "Rainbow Series" because each cover in the series is a different color.[3] Systems that followed the procedures of the series were considered *trusted* systems, due to the rigid requirements for following the configuration criteria contained within this series. The Rainbow Series documents have evolved into today's *Common Criteria Evaluation and Validation Scheme* (CCEVS)[4] to meet the demands of processing and computer systems of today; however, the fundamentals of how to secure and audit systems and information have remained semiconstant and have begun to be defined as best practices.

The term *best practices* is vague, but the concept comprises a set of standards, experiences, and after-action reports (lessons learned) that are intended to help meet the information security controls most commonly expected of trusted organizations. These best practices were derived from standards that were developed for secure or trusted systems. Today, information assurance best practices have evolved from the rigid requirements of the Rainbow Series standards. Best practices can now be found contained in various documents, including:

- Auditing checklists that can be obtained from the Information Systems Audit and Control Association (ISACA) as discussed in Section 54.4.3.3 in this chapter
- Defense Information Systems Agency's Security Technical Implementation Guides (DISA-STIG)[5]
- National Institute of Standards and Technology (NIST) Computer Security Division's Special Publications (SP) series—800 documents[6]

Another way to address the multiple regulations is to incorporate one of the more comprehensive frameworks, such as Control Objectives for Information and related Technology (COBIT, Section 54.4.3.3 in this chapter), Information Technology Infrastructure Library (ITIL),[7] and the International Organization for Standardization (ISO)[8] 27000 series (Section 54.3.2), with many relying on the others for ensuring that the information security areas have appropriate controls. Sections relevant to information security vary on the types of framework that should be incorporated but have value in that they identify areas that should have manual or technology controls pertinent to those areas.

For example, the Federal Information Security Management Act (FISMA) of 2002, the Sarbanes-Oxley Act of 2002, and the Health Insurance Portability and Accountability Act of 1996 all address security policy, logical access controls, and methods for the handling and storage of sensitive information.

In the past, if help was needed to classify information, the *Yellow Book*, or the CSC-STD-003-85, was the reference or standard. Once information was appropriately classified, an information security classification policy was developed by the organization and was applied to securing that data in information systems.

Today, legislative and industry specific requirements sometimes define what is critical or sensitive information; for many organizations looking for answers, the starting reference is the NIST SP 800-60, *Guide for Mapping Types of Information and Information Systems to Security Categories*. The NIST SP 800 series provide a wealth of information to support effective IT and security audits.

54.2.1 Gramm-Leach-Bliley Act. The Financial Services Modernization Act of 1999, more commonly known as the Gramm-Leach-Bliley Act (GLBA), regulates the security of consumers' personal financial information and nonpublic personal information (NPI). Major provisions of the GLBA address privacy requirements for information, disclosures of personal and financial information to third parties, rules regarding the safeguarding of information, and protections against the use of pretexting to illegally acquire personal information.[9] Financial information may include names of financial institution customers, account numbers, and information related to balances and transactions. NPI may include date of birth, Social Security number, driver's license number, financial account numbers, balances, home address, telephone numbers, and so on.

GLBA also governs the disclosure of such information to third parties. Institutions that are affected by GLBA must disclose to customers, in the form of a privacy notice, when and how their information is released. Additionally, institutions must provide for customers the opportunity to opt out of having their information shared with some third parties. Customers may not use the opt-out privileges to stop the release of their information to law enforcement agencies that are performing an authorized investigation.

GLBA regulations also protect consumers against the increasingly used form of social engineering called "pretexting." Pretexting is the act of using false statements or impersonation to obtain access to the information of another individual. An example would be a perpetrator calling a financial institution and pretending to be an authorized account holder requesting to make a funds transfer by telephone. In the absence of face-to-face verification of a drivers' license or some other token, banks are allowing such transactions to occur after the caller has verified some "known" information, be it date of birth, address, Social Security number, and so on. GLBA makes the act of pretexting illegal and encourages institutions to use more stringent methods for confirming a customer's identity.

Internal auditing to determine compliance with GLBA protects organizations and ensures the proper handling of NPI. Because mishandling or unauthorized release of NPI can lead to incidents of fraud and identity theft, institutions that are required to use and maintain such information must be very diligent in their security processes. Periodic audits of business and network processes that control the storage, transmission, and use of NPI will aid in alerting organizations to weaknesses in their security and will document due diligence efforts in the protection of customer information.

For more details of compliance with GLBA, see Chapter 64 in this *Handbook*.

54.2.2 Federal Information Systems Management Act. FISMA affects almost all U.S. Governmental entities and was created as part of the E-Government Act (Public Law 107-347)[10] passed by the 108th Congress in 2002 and signed into law by then President George W. Bush. The resolution recognized the importance of information security to the economic and national security interests of the United States. Title III of the E-Government Act entitled the FISMA[11] to require each federal agency to develop, document, and implement an agency-wide information assurance program that would provide security controls for the information and systems that support the operations and assets of the agency. This measure should also include security controls for information-handling services or systems that are provided, or managed, by third-party agencies, contractors, or other sources.

FISMA defined what should be included in an information-security program to be effective:

- Periodic assessments of risk, including the magnitude of harm that could result from the unauthorized access, use, disclosure, disruption, modification, or destruction of information and information systems that support the operations and assets of the organization (Many corporate entities may refer to this as a business impact analysis [BIA].)
- Policies and procedures that are based on the information provided by risk assessments, that would cost effectively reduce information security risks to an acceptable level, and that would ensure that information security is addressed throughout the life cycle of each organizational information system
- Subordinate plans for providing adequate information security for networks, facilities, information systems, or groups of information systems, as appropriate
- Security awareness training to inform personnel (including contractors and other users of information systems that support the operations and assets of the organization) of the information-security risks associated with their activities and their responsibilities in complying with organizational policies and procedures designed to reduce these risks
- Periodic testing and evaluation of the effectiveness of information-security policies, procedures, practices, and security controls will be performed, with frequency determined by the level of risk, but no less than annually
- A process for planning, implementing, evaluating, and documenting remedial actions to address any deficiencies in the information-security policies, procedures, and practices of the organization
- Procedures for detecting, reporting, and responding to security incidents
- Plans and procedures to ensure continuity of operations for information systems that support the operations and assets of the organization

FISMA, along with the Paperwork Reduction Act of 1995 and the Information Technology Management Reform Act of 1996 (Clinger-Cohen Act), explicitly emphasizes a risk-based policy for cost-effective security. In supporting and reinforcing this legislation, the Office of Management and Budget (OMB) through Circular A-130, Appendix III, Security of Federal Automated Information Resources,[12] requires executive agencies within the federal government to:

- Plan for security
- Ensure that appropriate officials are assigned security responsibility
- Periodically review the security controls in their information systems
- Authorize system processing prior to operations and periodically thereafter

These management requirements presume that responsible agency officials understand the risks and other factors that could adversely affect their objectives. Moreover, these officials must understand the status of their security programs, and the security controls planned or in place to protect their information and information systems, in order to make informed judgments and investments that appropriately mitigate risk to an acceptable level.

The ultimate objective is to conduct the day-to-day operations of the agency and to accomplish the agency's stated missions with security commensurate with level of risk, while acknowledging the magnitude of harm that could result from the unauthorized access, use, disclosure, disruption, modification, or destruction of information. As a key element of the FISMA Implementation Project, NIST also developed an integrated risk framework[13] that effectively brings together all of the FISMA-related security standards and guidance, in order to promote the development of comprehensive and balanced information-security programs by agencies.[14]

54.2.3 Health Insurance Portability and Accountability Act of 1996. The Health Insurance Portability and Accountability Act (HIPAA),[15] 42 USC 201 et seq., includes specific provisions affecting information security:

> The Office for Civil Rights enforces the HIPAA Privacy Rule, which protects the privacy of individually identifiable health information; the HIPAA Security Rule, which sets national standards for the security of electronic protected health information; the HIPAA Breach Notification Rule, which requires covered entities and business associates to provide notification following a breach of unsecured protected health information; and the confidentiality provisions of the Patient Safety Rule, which protect identifiable information being used to analyze patient safety events and improve patient safety.[16]

For extensive discussion of HIPAA, see Chapter 71 in this *Handbook*.

54.2.4 Payment Card Industry Data Security Standards. The Payment Card Industry Security Standards Council was formed in 2006 by collaboration of American Express, Discover Financial Services, JCB International, MasterCard Worldwide, and Visa Inc. The Council has promulgated the

- Payment Card Industry Data Security Standard (PCI DSS) for "merchants and entities that store, process, or transmit cardholder data"
- Payment Application Data Security Standard (PA-DSS) for software developers
- PIN Transaction Security (PTS) requirements for manufacturers of card-reading equipment.[17,18]

The core goals of the PCI DSS are summarized in the organization's table of goals and requirements as shown in Exhibit 54.1.[18]

The official site[19] has an extensive range of documents available at no cost (simple registration required), including (as of July 2013):

- FAQs for use with ROC [Report on Compliance] Reporting Instructions for PCI DSS v2.0
- Glossary v2.0
- Navigating the PCI DSS v2.0
- PCI DSS AOC—Merchants v2.0
- PCI DSS AOC—Service Providers v2.0
- PCI DSS Quick Reference Guide v2.0
- PCI DSS Summary of Changes Version 1.2.1 to 2.0
- PCI Quick Reference Order Form

EXHIBIT 54.1 Organization's Goals and Responsibilities

Goals	PCI DSS Requirements
Building and Maintain a Secure Network	**1.** Install and maintain a firewall configuration to protect cardholder data **2.** Do not use vendor-supplied defaults for system passwords and other security parameters
Protect Cardholder Data	**3.** Protect stored cardholder data **4.** Encrypt transmission of cardholder data across open public networks
Maintain a Vulnerability Management Program	**5.** Use and regularly update anti-virus software or program **6.** Develop and maintain secure systems and applications
Implement Strong Access Control Measures	**7.** Restrict access to cardholder data by business need to know **8.** Assign a unique ID to each person with computer access **9.** Restrict physical access to cardholder data
Regularly Monitor and Test Network	**10.** Track and monitor all access to network resources and cardholder data **11.** Regularly test security systems and processes
Maintain an Information Security Policy	**12.** Maintain a policy that addresses information security for all personnel

- Prioritized Approach for PCI DSS Version 2.0
- Prioritized Approach Tool Version 2.0
- ROC Reporting Instructions for PCI DSS v2.0
- Summary of Changes from Prioritized Approach for PCI DSS Version 1.2 to 2.0
- Summary of PCI DSS and PA-DSS Feedback 2012

The specific guidelines can be integrated into audit patterns adapted for the requirements of each organization handling credit card data.

54.2.5 Publicly Available Security Publications. The NIST SP (SP) 800-53 Revision 1 is the "Recommended Security Controls for Federal Information Systems."[20] The purpose of SP 800-53 is to provide guidelines for selecting and specifying what should be security controls for information systems supporting the executive agencies of the federal government. The guidelines have been developed to help achieve secure system requirements by:

- Facilitating a more consistent, comparable, and repeatable approach for selecting and specifying security controls for information systems

- Providing a recommendation for minimum security controls for information systems categorized in accordance with Federal Information Processing Standard (FIPS) 199,[21] *Standards for Security Categorization of Federal Information and Information Systems*
- Providing a stable yet flexible catalog of security controls for information systems to meet current organizational protection needs and the demands of future protection needs based on changing requirements and technologies
- Creating a foundation for the development of assessment methods and procedures for determining security control effectiveness

The guidelines provided in SP 800-53 are applicable to all federal information systems[22] with the exception of systems designated as national security systems as defined by 44 USC Section 3542.[23] The guidelines have been broadly developed from a technical perspective to complement similar guidelines for national security systems. This publication is intended to provide guidance to federal agencies implementing FIPS 200,[24] *Minimum Security Requirements for Federal Information and Information Systems*. In addition to the agencies of the federal government, the state, local, tribal governments, and private-sector organizations that comprise the critical infrastructure of the United States are also encouraged to use these guidelines as appropriate.

The security controls in SP 800-53 have been developed using inputs from a variety of sources including NIST SP 800-26, Department of Defense (DoD) Policy 8500, Director of Central Intelligence Directive (DCID) 6/3, ISO/IEC Standard 17799, General Accounting Office (GAO) Federal Information System Controls Audit Manual (FISCAM), and Health and Human Services (HHS) Centers for Medicare and Medicaid Services (CMS) Core Security Requirements. The security controls cover these topic areas[25]:

- Access control
- Accountability and audit
- Awareness and training
- Certification, accreditation, and security assessments
- Configuration management
- Contingency planning
- Identification and authentication
- Incident response
- Maintenance
- Media protection
- Personnel security
- Physical and environmental protection
- Risk assessment
- Security planning
- System and communications protection
- System and information integrity
- System services and acquisition

54.2.6 Risk Framework. The following list illustrates an example of how to create a security-audit framework that draws on a wide array of audit standards:

1. Categorize the information system and the information resident within that system based on impact. (FIPS 199iii and NIST SP 800-60[26])
2. Select an initial set of security controls for the information system based on the FIPS 199 security categorization and apply tailoring guidance, as appropriate, to obtain a starting point for required controls. (FIPS 200 and NIST SP 800-53)
3. Supplement the initial set of tailored security controls based on an assessment of risk and local conditions, including organization-specific security requirements, specific threat information, cost-benefit analyses, and special circumstances. (NIST SP 800-53 and SP 800-30[27])
4. Document the agreed-on set of security controls in the system security plan, including the organization's justification for any refinements or adjustments to the initial set of controls. (NIST SP 800-18[28])
5. Implement the security controls in the information system
6. Assess the security controls using appropriate methods and procedures to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system. (NIST SP 800-53A[29])
7. Authorize information system operation based on a determination of the risk to organizational operations, organizational assets, or to individuals resulting from the operation of the information system and the decision that this risk is acceptable. (NIST SP 800-37[30])
8. Monitor and assess selected security controls in the information system on a continuous basis including documenting changes to the system, conducting security impact analyses of the associated changes, and reporting the security status of the system to appropriate organizational officials on a regular basis.[31]

54.3 AUDITING STANDARDS. Auditing standards and practices assist organizations in meeting regulatory requirements and provide industry guidance for ensuring due diligence, certification, and ensuring stakeholder security. As discussed in the preceding section, standard-setting entities include international organizations, federal and state levels of government, and industry-specific groups or associations whose focus includes regulatory compliance. Guiding principles and information published by these entities may be required or, at the very least, provide managers with the newest information regarding industry advances. Managers must not only be aware of new developments in auditing methodology; they must also be adept in identifying and applying the best solution available to their unique physical and logical architecture.

54.3.1 Introduction to ISO. ISO[32] is a nongovernmental cooperative of international members who work together to create, identify, and publish industry standards as they relate to business and technology. Membership comprises member bodies, correspondent members, and subscriber members.

- A *member body* is a group that represents the accepted standards of a specific country.

- *Correspondent members* represent a country that does not yet have fully developed standards.
- *Subscriber members* represent countries with developing economies.[33]

Level of membership determines technical and policy participation and voting privileges, but members all are equally represented in international standardization and are encouraged to contribute.

Standards are developed based on industry interest; needs of vendors, governments, and consumers; and initiatives of the ISO where it sees market interest. Once the idea of a standard has been recognized by the ISO, member committees begin the process of defining the technical requirements necessary for the standard. Committee members include technical experts representing countries interested in the potential new standard. Once the technical requirements have been defined, interested members and countries work toward agreement on detailed specifications for the standard. When final details have met committee approval, a draft of the international standard is released for review. Final release of an official ISO standard[34] requires approval of two-thirds of the members who participated in the development process and 75 percent of all members who vote.[35]

Standards established by ISO are not regulator; they represent best-practice guidance that is globally consistent for a given trade or line of business. Widespread acceptance and implementation of ISO standards and guidance further elevates them as identifiers of distinction within a given industry. Similar to worldwide recognition of ISO 9000 and ISO 14000 as standards for Quality and for Environmental Management (respectively), the ISO 27000 series is quickly becoming accepted as the measurement of excellence in information-security management.

54.3.2 ISO/IEC 27000 Series. Information and network security has been a concern since the early days of worms and viruses; however, it has only been in the past 15 to 20 years that structures have been defined to integrate security with network use and risk management. Introduction of BS7799 and its acceptance by ISO have given administrators the long-needed tools to identify and manage these domains. Additionally, regulations such as the Sarbanes-Oxley Act of 2002 (SOX), the Gramm-Leach-Bliley Act (GLBA), Health Insurance Portability and Accountability Act (HIPAA), and the Federal Information Security Management Act (FISMA) have set down compliance requirements that are more easily achievable with a comprehensive information-security management system in place.

British Standard 7799, representing best practices for information security, was first published in February 1995. This three-part standard encompasses security objectives and controls as well as the development methodology necessary to create an information-security management system (ISMS).[36]

- BS7799: Part 1, Best Practices for Information Security Management
- BS7799: Part 2, Defines specifications for information-security management systems
- BS7799: Part 3, Guidelines for information-security risk management

In the past decade, British Standard 7799 was the basis for the ISO 17799 and then into the ISO 27000 series, which continues to evolve as a widely accepted standard for information-security management and for operational procedure guidance. Both

agencies (and their collective members) understood the need for a comprehensive program that addressed the risk management and auditing needs of organizations while defining a framework for programs and policies to support this concept. Partial or full implementation of these standards not only prepares organizations for audits and risk assessments but can also help them meet regulatory requirements.

The ISMS defined by BS7799 and further improved on by ISO 17799 and ISO/IEC 27000 is a means to ensure the best environment not only for security practices but also for objective auditing and assessments. An organization's interpretation of, and adherence to, the ISMS model can be a major indicator of the security climate of the organization as well as of management's support of policies and procedures. Programs that support ISMS functions must be promoted and used by management to demonstrate their significance to employees and to encourage participation. Audits typically consist of a planning phase, the actual audit process, and the final report to the audit committee presenting the details of audits or assessments.

In December 2000, BS7799 Part 1 (Best Practices for Information Security Management) was revised and released as ISO 17799.[37] The ISO 17799 standard comprises 10 domains:

1. Business continuity planning
2. System access control
3. System development and maintenance
4. Physical and environmental security
5. Compliance
6. Personnel security
7. Security organization
8. Computer and operations management
9. Asset classification and control
10. Security policy

In 2005, ISO completed a cooperative effort with the International Electrotechnical Commission to release ISO/IEC 17799:2005. This release, entitled *Information Technology—Security Techniques—Code of Practice for Information Security Management*, added objectives and controls for information-security management and updated the previous release to include advances in technology (mobile and wireless) and in applications and uses of the Internet.[38]

BS7799 and ISO 17799 established widely implemented guidance for the information-security management systems standard, but the evolutionary path of their published standards has made them nearly obsolete. The ISO/IEC 27000[39] set of standards not only incorporates the guidance of BS7799 and ISO 17799[40] but also provides a more thorough explanation and understanding of each domain by partitioning them into distinctly separate standards. The ISO/IEC ISMS family includes the following International Standards (as of July 2013):

- ISO/IEC 27000:2009, Information security management systems—Overview and vocabulary
- ISO/IEC 27001:2005, Information security management systems—Requirements
- ISO/IEC 27002:2005, Code of practice for information-security management

- ISO/IEC 27003:2010, Information-security management system implementation guidance
- ISO/IEC 27004:2009, Information-security management—Measurement
- ISO/IEC 27005:2011, Information-security risk management
- ISO/IEC 27006:2011, Requirements for bodies providing audit and certification of information-security management systems
- ISO/IEC 27007:2011, Guidelines for information-security management systems auditing
- ISO/IEC 27010:2012, Information-security management guidelines for intersector and interorganizational communications
- ISO 27799:2008, Health informatics—Information-security management in health using ISO/IEC 27002
- ISO/IEC FDIS 27013, Guidance on the integrated implementation of ISO/IEC 27001 and ISO/IEC 20000-1
- ISO/IEC TR 27008:2011, Guidelines for auditors on information-security management systems controls
- ISO/IEC TR 27015, Information-security management guidelines for financial services
- ISO/IEC WD 27016, Information-security management—Organizational economics
- ITU-T X.1051 | ISO/IEC 27011:2008, Information-security management guidelines for telecommunications organizations based on ISO/IEC 27002
- ITU-T X.1054 | ISO/IEC FDIS 27014, Governance of information security

Some ISO standards are freely available in electronic form from http://standards.iso.org/ittf/PubliclyAvailableStandards/index.html and they may also be purchased as electronic or printed documents from http://www.iso.org/iso/home/store.htm

ISO/IEC 27001 (ISMS Requirements) closely models the Organization for Economic Cooperation and Developments (OECD)[41] guidance on the security of information systems and networks, and is driven by the Plan-Do-Check-Act (PDCA) cycle.[42]

The PDCA cycle, developed by W. Edwards Deming in the 1950s, asks managers to view business processes as being in a constant feedback loop.[43] This method allows managers to be involved more thoroughly in the entire change control and improvement process.

Organizations may implement some or all of the standards of ISO/IEC 27001 or take their program a step further by applying for ISMS certification. Certification provides both auditors and stakeholders with the assurance that a comprehensive program is in place for information-security and business systems. Whether an organization seeks recognized ISMS certification or not, there are several advantages to adopting the objectives and controls outlined in ISO/IEC 27001. As stated earlier, having a comprehensive ISMS program in place can greatly contribute to compliance in other areas. An ISO/IEC 27001-certified ISMS will ensure compliance with the whole range of information-security–related legislation, including HIPAA, GLBA, and FISMA regulations, as well as the conditions for a successful Sarbanes-Oxley Section 404 report.[44]

Documented ISMS policies and programs, a security climate including both management and employees, and adherence to policy are not only necessary for an

information-security management system to be effective, but will ensure a successful, constructive, and effective audit process.

54.3.3 AICPA Security Audit Standards. Companies and government agencies are under constant pressure to improve efficiencies and to reduce operational costs. In addition, public companies are being driven by investors and shareholders to improve profitability, contributing in a small way to IT outsourcing in the 2000s. The U.S. Bureau of Labor Statistics began including "Cost control/cost cutting/increase profitability" as a reason for mass layoffs starting with the 2007 studies; the proportions of respondents using this entry as a reason for the layoffs were stable at approximately 3 percent of total yearly layoffs between 2007 and 2011.[45]

Outsourcing has not been limited to public companies; government agencies have also embraced outsourcing to contain costs. In 2013, data from the U.S. Federal Office of Management and Budget suggested that about one-third of all U.S. Government discretionary spending was outsourced.[46]

With the drive for more companies to focus on their core competencies, these types of business functions are being outsourced:

- Claims processing
- Customer service and helpdesk
- Data processing
- Financial and accounting
- Human resources and benefit plans
- Web hosting

Executive officers of organizations considering outsourcing must understand both the benefits and the risks; they also have to mitigate the risk to a level acceptable for their organizations.

The American Institute of Certified Public Accountants (AICPA) has defined three widely used security-audit standards, discussed in the following sections:

- SAS 70
- SSAE 16

ISAE 3402 is an international standard with which SSAE 16 complies.

54.3.3.1 Introduction to SSAE 16 Audits. The *Statement of Auditing Standards No. 70 (SAS 70): Service Organizations*, was originally created by the American Institute of Certified Public Accountants (AICPA) for financial-service reporting and was a key method of ensuring that a company being considering for supplying outsourced services met security requirements.[47]

The security-audit industry's need to report on security controls has led to the replacement of SAS 70 by the *Statement on Standards for Attestation Engagements No. 16 (SSAE 16), Reporting on Controls at a Service Organization*, which replaced the guidance for service auditors in SAS 70 for all such engagements ending on or after June 15, 2011.

EXHIBIT 54.2 Types of SSAE 16 Audits

SSAE 16 Report Content	Type I	Type II
Independent service audits report	Included	Included
Service organization's description of controls	Included	Included
Information provided by the independent service auditor; includes a description of the service auditor's tests of operating effectiveness and the results of those tests	Optional	Included
Other service relevant information	Optional	Included

SSAE 16 was also "drafted with the intention and purpose of updating the US service organization reporting standard so that it mirrors and complies with the new international service organization reporting standard—ISAE 3402."[48]

In addition to being used for evaluating outsourcing suppliers, the SSAE 16 can be used when *partnering* with companies that may have access to your IT resources and data—especially personally identifiable information (PII) about customers.

In this section, the two types of SSAE 16 audits and their components are described.

Terminology associated with SSAE 16 audits:

Service organization—the company that will process information or transactions on behalf of the user organization

Service auditor—retained by the service organization to conduct an SSAE 16 audit and provide a certification of audit results

User organization—the company that outsourced services and is the customer of the service organization

Users' auditors—either an internal or external auditing group that provides auditing services to the user organization

SOC—service organization controls; there are three types of SOC reports

The SSAE 16 defines the professional standards that are used by the service auditor, who will be assessing the internal controls that a service organization has implemented.

An SSAE 16 audit provides authoritative guidance that allows a service organization to disclose its control processes and activities to potential customers and to external auditors; the information is presented in a standard reporting format. There are three different types (SOC-1, -2, and -3) and two different levels of SSAE 16 audits (Type I and Type II), as illustrated in Exhibit 54.2.

- **SOC-1**: Controls at a service organization that are relevant to the user entity's internal controls over financial reporting; these can include both IT-related as well as financial controls. This report can be used by all types of service organizations; data centers may use this type of report in place of a SOC-2 if one or more financially related controls are reviewed.
- **SOC-2**: Strictly an IT-focused report with the controls preestablished and organized into the categories of security, availability, processing integrity, confidentiality, and privacy. The service organization may choose to have the audit focus on selected categories; all controls within the selected category/categories must be reviewed. Typically, these apply to data centers, ISPs, and ASPs.

- **SOC-3**: Similar to a SOC-2, this report focuses on trust services and principles; e.g., a Web trust audit. If the service organization elects to include the privacy principle in the audit scope, the criteria for the content of a privacy notice is included. This typically includes information about collecting the information, the choices that the individuals have regarding their personal information, the security of the information, and how the individuals can contact the service organization for inquiries, complaints, or disputes.

Each report contains several key sections:

- Opinion letter from the CPA firm.
- Management assertion regarding the description of the service organization's system and the suitability of the design of the controls; if the report is a Type II, then the assertion also includes the operating effectiveness of the controls in achieving the control objectives.
- Description of the controls and control environment.
- Results of the controls' testing.
- Key subservice organizations (e.g., outsourced IT monitoring) whose activities affect the controls.

A critical difference between an SSAE 16 Type I audit and Type II is that a Type I audit tests the control design effectiveness and tests one sample transaction for each control; this, to prove that the design of the control functions appropriately in at least one instance. For a Type II audit, the control design is tested as well as the controls' effectiveness over the audit period. Because of this testing, Type II audits are more costly but provide a detailed report and provide more value and reliance to the service organizations' users.

An SSAE 16 audit starts with the service auditor conducting an initial risk assessment and evaluation of the processing system and of the transaction systems processing controls, and a review of the policies and procedures that govern IT system and processes. Next, the auditor will develop a statement of work (SOW) that will be presented to the service organization with completion date, details of the audit, audit scope, and cost.

An SSAE 16 audit team usually will be managed by a CPA; a CPA must review the audit report and work and must sign the opinion. This individual will be the primary interface with the customer executive management and will certify the audit results. The audit team performing the work may be a single person or a team of perhaps two to four auditors. The technical lead is responsible for evaluation and testing of systems and network devices.

If a SOC-1 audit is being performed, there may be two to five IT-related control objectives, and within each control objective, there may be one to fifteen controls tested. Typically, there is more granularity in the testing of the controls, so that if a particular control fails (e.g., not disabling a terminated user's access promptly), then it would not cause other controls bundled together to fail (e.g., if user provisioning—add, change, or remove users—is listed as a single control instead of three controls). For an SOC-1 audit, the service organization is responsible for identifying and approving the control objectives and the controls; the service auditor often provides information as to best practices and controls that typically are tested. For the SOC-2 and SOC-3 reports, the controls are preestablished, and the service organization may select which of the key categories are to be reviewed in the audit.

In addition to reviewing hardware and software configurations, IT policies and procedures typically are included to be evaluated, reviewed, and verified. Management, including system and network administrators, will be interviewed in the walkthrough phase to determine how certain tasks are performed; a sample item will be tested at that time to verify that the control's design is effective (i.e., that the control works for a sample transaction). For example, who approves a user account request form, and how does the form flow through the approval process? System or application administrators will be asked to explain what they do from the time they receive an account request form until the account is created. The auditor will use the information gathered from the interviews to determine if account policies and procedures are being followed and if they are adequate for the customer's environment.

A brief overview of what an auditor would review for a user request policy and how the results would be reported follows:

- **Policy statement**. All users who require a new computer account or modification to an existing account, on any system at the company will be required to complete User Account Request Form UA-09.
- **Policies and procedures**. User Account Request Policy "IT-23 November 1, 2011." All requests for a new account or for modification of an existing account must be submitted on a UA-09 Form.
- **Test procedure**. Conduct a sample review of user request forms and review the forms for proper authorized signatures and completeness. Interview authorized signers for compliance with user request policy; ensure that users are only given roles consistent with their job position's least privileges.
- **Test results**. Collect statistical samples of user request forms and interview the supervisors of the individuals making the requests to determine that least privileges needed to perform their job have been requested and approved. Check forms for completeness and accuracy of information.
- **Analysis of test results**. The auditor will provide both the positive test results and negative results with commentary. A sample of a positive or negative test results:
  - **Positive test result**. All forms that were reviewed were completed in accordance with User Account Request Policy "IT-23 November 1, 2011," and all interviews indicated that all policies and procedures were correctly followed. No discrepancies were found.
  - **Negative test result**. After reviewing the User Account Request Form and interviewing supervisors and authorized signers, it was determined that User Account Request Policy "IT-23 November 1, 2011" was not followed. Users were given privileges beyond what was required to perform their job.

Exceptions and findings are reported in the final SSAE 16 report sent to management in these ways:

- If the auditor deems that it is relevant, an exception is reported, including a description of the exception.
- The auditor will consider any mitigating controls or any management responses to the exception.
- If, in the auditor's opinion, the finding has an impact on the control objectives, then it would be reported with a qualified opinion and potential impact; if not, it would be reported as an exception.

The final document that a SSAE 16 will provide is the conclusion and summary report, which will clearly state whether the description of the controls is accurately and fairly presented. The report will also state whether the controls were adequately designed to meet their control objectives. The report for a Type II audit would also include the results of tests that were performed and would evaluate whether they provide reasonable assurance that controls were operating correctly during a period of time, normally no less than six months and no more than 12 months; however, for the first SSAE 16, a service organization must have a minimum review period of three months.

54.3.3.2 Cost and Benefits of SSAE 16 Audits. Before undertaking an SSAE 16 audit, any organization, whether a service provider or an outsourcing company, must decide whether such a procedure is worth the costs. For example, a small service organization whose customer requests an SSAE 16 audit must determine whether to keep that customer: an initial audit Type II SSAE 16 audit's cost can range from $15,000 to over $1 million, in addition to the annual cost to keep the audit current. Because of the reduced testing, a Type I may be engaged for $10,000 to $13,000, depending on the number of controls, the control objectives, and the extent that the service auditor (SSAE 16 auditor) can rely on the service organization's internal audit work.

User organizations have quickly embraced the SSAE 16 audit because of the benefits in complying with the SOX and GLBA requirements. Depending on the scope and timing of the audit, regulators and/or financial auditors may be able to place reliance on the audit and reduce their own audit scopes; this can translate into time and expense savings.

The audit can help foster business development activities, since an SSAE 16 audit is frequently regarded as an affirmation of the service organization's commitment to quality and security; the audit report can also provide benefits and reduce paperwork when responding to requests for proposal (RFPs) or similar offerings.

SSAE 16 has adopted the *Committee of Sponsoring Organizations of the Treadway Commission* (COSO) as its standard for reviewing internal controls.[49] The COSO definition of internal control is "a process, effected by an entity's board of directors, management and other personnel, designed to provide reasonable assurance regarding the achievement of objectives." Also, COSO is the same standard that Sarbanes-Oxley Section 404 has adopted for how to review internal controls. See Section 54.5.3.2 for further details of COSO.

The SSAE 16 has pros and cons for both the service organization and user organization; these issues must be understood so that management can perform a cost-benefit analysis (see Exhibit 54.3).

54.3.3.3 SSAE 16 Audits Conclusion. In conclusion, an SSAE 16 audit is not a 100 percent guarantee of assurance that all security-related controls and processes have been correctly designed and implemented; after all, the service organization selects the controls to be tested in an SOC-1 engagement. For SOC-2 or SOC-3 engagements, high-value control objectives and controls are included; however, they do not cover all possible controls. Nonetheless, an SSAE 16 audit provides a high-level of confidence that risks have been assessed and mitigated.

The SSAE 16 audit results can also provide information about areas for further investigation. A major benefit of an SSAE 16 audit is that it can be used to meet the reporting requirements for SOX Section 404 reporting, help in meeting GLBA and

EXHIBIT 54.3 Pros and Cons for Service and User Organizations

Factor	Service Organization	User Organization
Independent assessment of controls	Pro	Pro
Lower cost for evaluation of controls	Con	Pro
No additional review of controls required	Con	Con
SSAE 16 audits are forward looking	Con	Con
SSAE 16 audits must be continuously reviewed and updated	Con	Pro
SSAE 16 audits increase the value of services	Pro	Pro
Disruption to service organization is reduced by eliminating user organization auditors from visiting service organization to conduct audit	Pro	Pro
The SSAE 16 audit can be used to build a strong working relationship between organizations	Pro	Pro
Audit results can provide opportunities for improvements	Pro	Pro

HIPAA compliance, and may provide benefits as a marketing tool regarding service quality.

54.3.4 ISAE 3402 Standards. The *International Standard on Assurance Engagements (ISAE) No. 3402, Assurance Reports on Controls at a Service Organization*, is the international standard under which SSAE 16 has been approved:

> [ISAE 3402] was issued in December 2009 by the International Auditing and Assurance Standards Board (IAASB), which is part of the International Federation of Accountants (IFAC). ISAE 3402 was developed to provide an international assurance standard for allowing public accountants to issue a report for use by user organizations and their auditors (user auditors) on the controls at a service organization that are likely to impact or be a part of the user organization's system of internal control over financial reporting.[50]

Ernst & Young describe ISAE 3402 as follows in a 2013 brochure:

> Although many businesses have been outsourcing portions of their work for years now, outsourcing is still becoming more popular by the day. This is especially driven by increased globalization, technological evolutions, and the need for standardised business processes. Outsourcing is any task, operation, job, or process that could be performed by employees within the user organisation, but is instead contracted to a third party (service organisation) or another group company for a period of time.
>
> Some examples for the financial sector are:
>
> - Asset managers that perform asset management services for different parties within the group company.
> - Pension administrators who perform the administration for pension funds.
> - Claim service companies that perform claim handling services for large insurers.
>
> The widespread use of outsourcing requires organizations to better manage their risks associated with the outsourced services. More specifically, the user organisation requires a degree of assurance that the service organisation has a well-established internal control framework that is operating effectively. New regulations, regulatory authorities and supervisory boards also ask for specific controls over outsourced procedures.[51]

54.4 SARBANES-OXLEY ACT

54.4.1 Introduction. This section focuses on the practicalities of SOX compliance in the real world and introduces concepts that the practitioner is sure to encounter.

Sarbanes-Oxley is the financial reporting act enacted into U.S. law in July 2002.[52] It was the U.S. Government's response to a number of major public-company accounting failures, including Enron[53] and WorldCom.[54]

SOX is designed to reduce fraud and conflicts of interest while increasing financial transparency and public confidence in the market and in companies' internal control systems. The act is intended to encourage honesty and integrity in financial reporting and to place responsibility directly with owners, managers, and executives. The act requires that the executive officers certify that they are "responsible for establishing and maintaining internal controls" and "have designed such internal controls to ensure that material information relating to the company and its consolidated subsidiaries is made known to such officers by others within those entities, particularly during the period in which the periodic reports are being prepared." SOX allows for the possibility of both *civil* litigation and *criminal* prosecution.[55] Thus, noncompliance is likely to result in fines, imprisonment, or both for those who fail to ensure the accuracy of their organization's financial reporting.

Bearing in mind that SOX currently applies only to *publicly traded* companies with a market capitalization (or public debt) *exceeding $75 million*, many companies delisted themselves from U.S. stock exchanges or reverted to private status to avoid inclusion in the act.

More significantly, many non-U.S.-held companies now find themselves SOX-applicable because they or their subsidiaries are listed on the U.S. stock exchanges.[56] Compliance is necessary for the entire organization, not just for the listed subsidiary. Thus, from humble beginnings, the act now has global market penetration.

SOX may originally have been the driver for complying with stricture regulations; however, in the decade since it was passed, "SOX compliance has moved beyond being just a regulatory obligation. Today it is viewed in the context of a larger Governance, Risk, and Compliance (GRC) program that is closely integrated with strategic decision-making, aligned with business goals, and enabled by technology."[57]

It is in the best interests of an organization to design and implement a comprehensive governance program, to assess and implement internal processes and controls, and to set targets beyond that of mere regulatory compliance. The entire compliance effort should be repeatable, as compliance is not a one-off activity; rather it will be repeated annually for the life of the organization (or until such time as the law is materially revised or repealed).

54.4.2 Section 404. SOX consist of a number of sections; most are beyond the scope of this chapter. From an information assurance professional's viewpoint however, Section 404 is the most significant and most widely discussed. It is the area that directly addresses the impact of information technology on financial reporting and the adequacy and effectiveness of associated internal controls.

Section 404 reads as follows:

> SEC. 404. MANAGEMENT <<NOTE: 15 USC 7262.>> ASSESSMENT OF INTERNAL CONTROLS.
>
> **(a)** Rules Required.–The Commission shall prescribe rules requiring each annual report required by section 13(a) or 15(d) of the Securities Exchange Act of 1934 (15 USC 78m or 78o(d)) to contain an internal control report, which shall—

> **(1)** state the responsibility of management for establishing and maintaining an adequate internal control structure and procedures for financial reporting; and
>
> **(2)** contain an assessment, as of the end of the most recent fiscal year of the issuer, of the effectiveness of the internal control structure and procedures of the issuer for financial reporting.
>
> **(b)** Internal Control Evaluation and Reporting.–With respect to the internal control assessment required by subsection (a), each registered public accounting firm that prepares or issues the audit report for the issuer shall attest to, and report on, the assessment made by the management of the issuer. An attestation made under this subsection shall be made in accordance with standards for attestation engagements issued or adopted by the Board. Any such attestation shall not be the subject of a separate engagement.

With widespread reliance on IT for financial and operational management systems, it is now accepted within U.S. industry that adequate and effective internal controls are necessary, particularly for systems supporting significant business processes.[58]

54.4.3 Achieving Compliance. SOX compliance should be approached by identifying key processes within the organization, documenting how they are implemented and controlled, and the defining or characterizing methods for reporting on their success or failure.

Coverage must include the entire system life cycle, and internal processes must include projects, design, architecture, system development, delivery, and operations. An auditor will be looking for no more than the assessment of core processes, the adequacy of associated controls, and whether those processes were followed, and controls executed, as stated. For more on SOX issues and regulations, see Chapter 64 in this *Handbook*.

54.4.3.1 Control Framework. Fortunately, the hard work of choosing a controls and assessment framework has already been done. The U.S. Securities and Exchange Commission (SEC) mandates the use of the internal control framework recommended by COSO,[59] and the Public Company Accounting and Oversight Board (PCAOB)[60] supports the SEC's choice as detailed in its *Auditing Standard No. 5, An Audit of Internal Control Over Financial Reporting Performed in Conjunction with an Audit of Financial Statements*.[61]

54.4.3.2 COSO. The COSO "Internal Control—Integrated Framework (2013)"[62] was sponsored by a consortium comprising:

- American Accounting Association (AAA)
- American Institute of Certified Public Accountants (AICPA)
- Financial Executives International (FEI)
- Institute of Internal Auditors (IIA)
- Institute of Management Accountants (IMA)

The Executive Summary discusses changes in the framework since the original 1992 version, including:

- Expectations for governance oversight
- Globalization of markets and operations
- Changes and greater complexities of business
- Demands and complexities in laws, rules, regulations, and standards
- Expectations for competencies and accountabilities

- Use of, and reliance on, evolving technologies
- Expectations relating to preventing and detecting fraud.

According to COSO, "Internal control is a process, effected by an entity's board of directors, management, and other personnel, designed to provide reasonable assurance regarding the achievement of objectives relating to operations, reporting, and compliance."

The Executive Summary continues:

> This definition reflects certain fundamental concepts. Internal control is:
>
> - *Geared to the achievement of objectives* in one or more categories—operations, reporting, and compliance
> - *A process* consisting of ongoing tasks and activities—a means to an end, not an end in itself
> - *Effected by people*—not merely about policy and procedure manuals, systems, and forms, but about people and the actions they take at every level of an organization to affect internal control
> - Able to *provide reasonable assurance*—but not absolute assurance, to an entity's senior management and board of directors
> - *Adaptable to the entity structure*—flexible in application for the entire entity or for a particular subsidiary, division, operating unit, or business process.

The "five integrated components" of internal control in the COSO framework can be summarized as follows:

- *Control environment:* "… [T]he set of standards, processes, and structures that provide the basis for carrying out internal controls across the organization. …"
- *Risk assessment:* Evaluation of probabilities for identified harmful events in the business context for the organization.
- *Control activities:* "[P]reventive or detective" actions taken to mitigate risks while achieving the corporate objectives.
- *Information and communication:* Organizations need current, accurate data about all aspects of operations, including monitoring for internal controls. "Communication is the continual, iterative process of providing, sharing, and obtaining necessary information."
- *Monitoring activities:* "Ongoing evaluations … provide timely information. Separate evaluations, conducted periodically, will vary in scope and frequency depending on assessment of risks, effectiveness of ongoing evaluations, and other management considerations. Findings are evaluated against criteria established by regulators, recognized standard-setting bodies or management and the board of directors, and deficiencies are communicated to management and the board of directors as appropriate."

54.4.3.3 COBIT. ISACA[63] was originally called the *Information Systems Audit and Control Association*; ISACA defined a practical framework from which to derive key processes and controls for SOX compliance activities using its COBIT (Control Objectives for Information and Related Technology) framework. ISACA describes the principles underlying the framework as follows:

> COBIT 5 is based on five key principles for governance and management of enterprise IT:
>
> - Principle 1: Meeting Stakeholder Needs
> - Principle 2: Covering the Enterprise End-to-End

- Principle 3: Applying a Single, Integrated Framework
- Principle 4: Enabling a Holistic Approach
- Principle 5: Separating Governance from Management

ISACA describes the benefits of COBIT 5 as follows:

COBIT 5 helps enterprises of all sizes:

- Maintain high-quality information to support business decisions
- Achieve strategic goals and realize business benefits through the effective and innovative use of IT
- Achieve operational excellence through reliable, efficient application of technology
- Maintain IT-related risk at an acceptable level
- Optimize the cost of IT services and technology
- Support compliance with relevant laws, regulations, contractual agreements and policies

The COBIT 5 framework describes seven categories of enablers:

1. Principles, policies, and frameworks are the vehicle to translate the desired behavior into practical guidance for day-to-day management.
2. Processes describe an organized set of practices and activities to achieve certain objectives and produce a set of outputs in support of achieving overall IT-related goals.
3. Organizational structures are the key decision-making entities in an enterprise.
4. Culture, ethics, and behavior of individuals and of the enterprise are very often underestimated as a success factor in governance and management activities.
5. Information is required for keeping the organization running and well governed, but at the operational level, information is very often the key product of the enterprise itself.
6. Services, infrastructure, and applications include the infrastructure, technology, and applications that provide the enterprise with information technology processing and services.
7. People, skills, and competencies are required for successful completion of all activities, and for making correct decisions and taking corrective actions.[64]

COBIT is a comprehensive framework for management of the governance of risk and control of IT, comprising four domains, 34 IT processes, and 215 control objectives.[65]

Although it has a wider purpose than simply addressing SOX compliance, COBIT removes much of the guesswork by suggesting 12 out of the 34 available standard processes for inclusion in a SOX implementation:

1. Acquire and maintain application software. (AI2) The provision of line of business applications, whether custom or off the shelf, should be congruent with business requirements.
2. Acquire and maintain technology infrastructure. (AI3) Technology infrastructure should underpin the delivery of services required by the business and follow organizational standards.
3. Enable operations. (AI4) Support effective use of technology through knowledge transfer, agreed service levels, and properly integrated solutions.

4. Install and accredit solutions and changes. (AI7) Ensure that supplied systems are fit-for-purpose prior to delivery and that they function as expected post-installation.
5. Manage changes. (AI6) Provide an effective, formal change management system to reduce errors and defects that may impact successful operations and business processes.
6. Define and manage service levels. (DS1) Align key business processes with delivery of technology in a set of meaningful service definition documents, with a mechanism for reporting to stakeholders on performance and achievement of service levels.
7. Manage third-party services. (DS2) Ensure the delivery of essential third-party services to support business activities while remaining transparent with regard to cost, benefits, and risks.
8. Ensure systems security. (DS5) Maximize the integrity and availability of systems and the confidentiality of data stored thereon while minimizing the impact of vulnerabilities through the application of effective security policies, standards, and procedures.
9. Manage the configuration. (DS9) Maintain an accurate and readily accessible repository of application and system configuration information, baseline standards, and variations thereof.
10. Manage problems and incidents. (DS8, DS10) Provide formal problem and incident management for all operations, including identification of issues, root cause analysis, and end-to-end problem resolution.
11. Manage data. (DS11) Provide tried and tested backup, recovery, and disposal systems.
12. Manage the physical environment and operations. (DS12) Control access to facilities and to computing assets, to reduce the likelihood of business disruption resulting from data leakage or theft of assets.

Specific analysis may require the addition of further processes, and within each process, the addition or removal of specific controls to further tailor the framework for each organization. Process and control selection must be agreed on in collaboration with the internal-audit function and with external auditors before committing to a particular framework.

54.4.3.4 Testing. Of equal importance to process and control selection is the agreement of the testing plan and schedule, especially the number of tests and the associated sample sizes.

For most organizations, resources are not unlimited. It is likely that staff will be asked to assume responsibility for some portion of SOX activity in addition to their regular workload. Thus, achieving a balance between compliance and the number of controls associated with each process is critical, because it is the controls that will be tested throughout the year, with the test results supporting the SOX certification process. A reduction in the number of tests or sample sizes can directly affect the integrity and quality of the results and positively affect staff efficiency and productivity. Although testing will likely be seen as onerous, much of the control activity should be considered business as usual, and thus the day-to-day processes should yield much of the evidence required for testing.

This section focuses on the importance and impact of IT on the wider goal of SOX compliance, but not all processes, even within an IT organization, are technical. Many are activities such as director-level meetings, incident boards, recovery forums, project risk meetings, and so on, and all may be part of a process that is included under SOX and must therefore be tested.

Testing typically falls into two categories: examination and discussion. Examination requires that a representative sample of evidence showing that a control is operating properly be collated and verified as accurate. Discussion requires that understanding of the process and associated controls be confirmed through one-to-one or group meetings among responsible staff. Sample sizes will depend on the test population and the frequency of sampling, and probably on the importance of the desired degree of accuracy and precision of the results for statistical confidence levels.

54.4.4 Audit and Certification. There will be at least two independent audit processes at work: internal and external. Internal audit will occur at the culmination of SOX testing, as a final quality-assurance checkpoint that the organization is compliant. Any identified gaps at must either be remediated or officially acknowledged within a formal risk process prior to the arrival of the external auditors.

54.4.4.1 External Audit. External audit will occur close to the end of the year and is the final test of compliance. It is not acceptable to wait until the external audit function identifies a gap before carrying out remediation activities. Any areas of noncompliance at this stage will be officially reported and may be included in the final management letter.

54.4.4.2 Schedule. Certification requirements vary by organization, although obviously at least an annual certificate will be required. Many organizations certify quarterly or monthly as part of some other SOX-related or controls process, and use those certificates to supplement their final annual declaration (see Exhibit 54.4).

54.4.5 Sarbanes-Oxley Conclusion. SOX compliance activities should not stand alone within the organization. They should be considered part of a wider risk-management program, preferably one supported by an embedded, proactive, risk-aware control culture. Further, it is never recommended to execute a compliance program with the goal of simply meeting regulatory requirements. In reality, a program that is designed from the outset to exceed the current regulation, and be adaptable to changes

EXHIBIT 54.4 Typical Annual SOX Schedule

Period*	Activity
January/February	Create and agree on controls and associated tests, certification schedule, and testing plan based on agreed processes
March to August	Execute testing (typically by process)
September	Retest where necessary
October	Internal audit
November/December	External audit, produce certificates, publish management letter

*The actual dates for compliance activities will vary with each organization and its respective financial years and filing deadlines.

to existing or entirely new regulation, represents a much smarter investment of time and effort.

For additional discussion of SOX compliance, see Chapter 64 in this *Handbook.*

54.5 TECHNICAL FRAMEWORKS FOR IT AUDITS. The best IT auditors struggle with the reality of acceptable risk and vague controls, looking for solid frameworks to help with their audit. Great research and a good framework can help determine the questions for many of the unique or new audits that have not been faced in the past. Three useful models, as discussed next, can be applied from a high-level perspective to make sure that all of the relevant questions have been covered. In addition to models, general best practices help round out the technical bases to make sure the audit reveals any material weaknesses in key controls.

54.5.1 Framework 1: People, Processes, Tools, and Measures. The first useful framework—People, Processes, Tools, and Measures (PPTM)—is simple conceptually, but its wide scope digs quickly into often-overlooked critical processes and important controls. Take a randomly picked application about which nothing is known, and apply PPTM concepts to the application to understand it and to develop a series of controls to test, and useful information to gather.

1. **People** build the application, input data into the application, and use the output or storage of the data handled by the application. People are a component that touches or uses the application, and as such should have controls that govern the interaction. For example, how are the people authenticated or authorized as administrators or users of the application?
2. **Processes** surround the operations of the application. What processes govern standard operating practices, and how can they be validated?
3. **Tools** are the physical controls, such as other applications that might interface with the application or ensure the quality or validation of the application. Other types of tool controls include checking the platform that the application uses. What physical controls existed during the implementation or are currently in operations to make sure the application is in a secure environment?
4. **Measures** are often-overlooked metrics, although they represent some of the most useful data in thorough audits. They can verify that the application is performing its function in support of its assigned business role. For example, metrics can report on the operating efficiency of an application and how well it might scale to meet customer needs over time.

54.5.2 Framework 2: STRIDE. STRIDE, an acronym based on the ordered list of issues shown below, encompasses six areas of risk common to different technologies. Much like PPTM, STRIDE will help in understanding how the area of risk might play a role in the technology to be audited and then in finding the controls that alleviate that risk. The keepers of the Open Web Application Security Project (OWASP, at www.owasp.com) provide an excellent example of applying STRIDE to Web applications.

1. **S—Spoofing** describes the risk of any user assuming the identity of another user or masking directly or indirectly the attributes of another user. For example, a user logging onto a banking Website should generally not be able to assume another user's identity.

2. **T—Tampering** with data describes the need for input validation, process validation, and output validation. The integrity of the data should be preserved at all times, and in some cases protected from tampering, with encryption or by independent verification.
3. **R—Repudiation** describes a scenario where a user can suggest that a transaction never occurred. Poor audit trails lead to this situation. Nonrepudiation is the control that protects against repudiation. It includes signing a message with date and time as a unique, nonduplicatable moment.
4. **I—Information** disclosure describes the additional controls to prevent information from leaving the system to unauthorized parties or programs.
5. **D—Denial** of service describes a condition where the application or technology could be maliciously or inadvertently removed from service. If this is a concern, controls should be in place such as redundant systems, backups, and measures to prevent denial-of-service events.
6. **E—Elevation** of privilege describes the situation where a user maliciously or inadvertently gains administrative rights.

For a more complete model of the elements of computer security incidents, see Chapter 8 in this *Handbook.*

54.5.3 Framework 3: PDIO. PDIO—*Plan, Design, Implement, and Operations*—borrows from Cisco System's project management processes and describes the life cycle of a technology.[66] This framework can be very powerful, especially when combined with PPTM.

Administrators address people, processes, tools, and measures at each step in a project's planning, designing, implementation, and operations. PDIO can be viewed from a macroperspective to understand an entire new project or from a microperspective to understand each smaller element. The big-picture understanding of an audit is the mark of a great result that provides incredibly valuable—and useful—information to customers so that they can strengthen their controls.

54.5.4 General Best Practices. This short list of common best practices should help with additional controls for testing. This is especially helpful during an audit of a new technology that does not yet have an appropriate audit program.

1. Apply defense in depth so that you have layered controls.
2. Use a positive security model. By default, deny actions and allow only a trusted or benign subset of available options. Negative security models trust everything by default and attempt to discover when to block offending behavior.
3. Fail safely so that people, equipment, data, and business processes are safe from harm.
4. Run with least privilege to prevent users from escalating their privileges above what is necessary to do their job.
5. Avoid security by obscurity. Try to hide data or options. Security by obscurity is setting up the technology for compromise.
6. Keep security simple. Overly complex systems are difficult to administer and often end up not secured.

7. Detect intrusions and keep logs. There should be an audit trail for everything important.
8. Never trust infrastructure and services, especially those coming from a third-party managed solution, where there might be different security policies and practices from the originating organization.
9. Establish secure defaults. Default to safe values and actions over defaulting to less secure values and actions. Another common mistake is not to specify an important value, causing issues later for the technology or application.
10. Use open standards because they are tested more thoroughly than proprietary technologies. Encryption in Digital Rights Management is an example where several purportedly secure private ciphers or cryptographic systems have been broken.

54.5.5 Technical Frameworks Conclusion. Frameworks can be valuable assets when approaching a technical audit, allowing quick determination of those areas where the audit is going to be most demanding and where compliance can be most effective.

54.6 FURTHER READING

Brotby, W. K., and G. Hinson. *PRAGMATIC Security Metrics: Applying Metametrics to Information Security*. Auerbach, 2013.

Cascarino, R. E. *Auditor's Guide to IT Auditing*, 2nd ed. Wiley, 2012.

Chorafas, D. *IT Auditing and Sarbanes-Oxley Compliance: Key Strategies for Business Improvement*. Auerbach, 2008.

Cox, F. D. *Information Security: Risk Management of GLBA Privacy and Service Provider Oversight*. Amazon Digital Services, 2011.

Davis, C., M. Schiller, and K. Wheeler. *IT Auditing Using Controls to Protect Information Assets*, 2nd ed. McGraw-Hill Osborne Media, 2011.

Davis, R. E. *IT Auditing: Assuring Information Assets Protection*. Amazon Digital Services, 2011.

Halpert, B. *Auditing Cloud Computing: A Security and Privacy Guide*. Wiley, 2011.

Hawker, A. *Security and Control in Information Systems: A Guide for Business and Accounting*. Routledge, 2000.

Hingarh, V., and A. Ahmed. *Understanding and Conducting Information Systems Auditing*. Wiley, 2013.

Moeller, R. R. *IT Audit, Control, and Security*. Wiley, 2012.

Moeller, R. R. *Executive's Guide to IT Governance: Improving Systems Processes with Service Management, COBIT, and ITIL*. Wiley, 2013.

Roebuck, K. *Information Security Audit: High-Impact Strategies—What You Need to Know: Definitions, Adoptions, Impact, Benefits, Maturity, Vendors*. Lightning Source, 2011.

Senft, S., F. Gallegos, and A. Davis. *Information Technology Control and Audit*, 4th ed. Auerbach, 2012.

Weiss, M., and M. G. Solomon. *Auditing IT Infrastructures for Compliance*. Jones & Bartlett Learning, 2010.

54.7 NOTES

1. S. Abelow, "Five Ways that an IT Audit Adds Value." Yahoo! Voices, January 13, 2011, http://voices.yahoo.com/five-ways-audit-adds-value-7567404.html
2. See, for example, "Guide for Assessing the Security Controls in Federal Information Systems and Organizations" NIST Special Publication SP 800-53A Rev. 1 (June 2010), http://csrc.nist.gov/groups/SMA/fisma/assessment.html
3. U.S. Department of Defence CSC, "Rainbow Series" (1980), available at www.fas.org/irp/nsa/rainbow.htm
4. National Information Assurance Partnership (NIAP), www.niap-ccevs.org
5. Information Assurance Support Environment (IASE), *STIGs: Security Technical Implementation Guides*, U.S. Department of Defense (DoD) Defense Information Systems Agency (DISA), http://iase.disa.mil/stigs
6. Computer Security Resource Center (CSRC) Special Publications (800 Series), National Institute of Standards and Technology (NIST) Information Technology Laboratory (ITL), http://csrc.nist.gov/publications/PubsSPs.html
7. Information Technology Infrastructure Library (ITIL), www.itil-officialsite.com
8. ISO is not an abbreviation of International Organization for Standardization but instead a shortened form of the Greek *isos*, meaning equal. The term was chosen to be understood in all translations and speak to the organization's intent of equal representation for all members regardless of their national economic status or country's level of development.
9. U.S. Senate Committee on Banking, Housing and Urban Affairs, Financial Services Modernization Act (1999), http://banking.senate.gov/conf/grmleach.htm
10. NIST, "Federal Information Security Management Act (FISMA) Implementation Project," http://csrc.nist.gov/groups/SMA/fisma/index.html
11. Federal Information Security Management Act of 2002 (FISMA), 44 USC 3541, *et seq*., www.law.cornell.edu/uscode/text/44/3541
12. OMB, Circular A-130—Revised, Transmittal Memorandum No. 4, "Memorandum for Heads of Executive Departments and Agencies; Subject: Management of Federal Information Resources," www.whitehouse.gov/omb/circulars/a130/a130trans4.html
13. NIST, "FISMA Risk Management Framework," http://csrc.nist.gov/groups/SMA/fisma/framework.html
14. NIST, "FISMA Detailed Overview," http://csrc.nist.gov/groups/SMA/fisma/overview.html
15. Health Information Privacy, U.S. Department of Health & Human Services (HHS), www.hhs.gov/ocr/privacy
16. HIPAA, www.gpo.gov/fdsys/pkg/PLAW-104publ191/html/PLAW-104publ191.htm
17. PCI Security Standards Council—About Us, https://www.pcisecuritystandards.org/organization_info/index.php
18. "PCI DSS Quick Reference Guide," https://www.pcisecuritystandards.org/documents/PCI%20SSC%20Quick%20Reference%20Guide.pdf
19. PCI Security Standards Council—Documents Library, https://www.pcisecuritystandards.org/security_standards/documents.php?document=pci_dss_v2-0#pci_dss_v2-0

20. NIST, "Recommended Security Controls for Federal Information Systems," Special Publication 800-53 Revision 4 (April 2013), http://dx.doi.org/10.6028/NIST.SP.800-53r4
21. NIST, "Standards for Security Categorization of Federal Information and Information Systems (February 2004)," FIPS 199, http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf
22. A *federal information system* is an information system used or operated by an executive agency, by a contractor of an executive agency, or by another organization on behalf of an executive agency.
23. NIST, "Guideline for Identifying an Information System as a National Security System (August 2004)," Special Publication 800-59, http://csrc.nist.gov/publications/nistpubs/800-59/SP800-59.pdf
24. NIST, "Minimum Security Requirements for Federal Information and Information Systems (March 2006)," FIPS-200, http://csrc.nist.gov/publications/fips/fips200/FIPS-200-final-march.pdf
25. NIST, "Federal Information Security Management Act (FISMA) Implementation Project," http://csrc.nist.gov/groups/SMA/fisma/index.html
26. NIST, "Draft Guide for Mapping Types of Information and Information Systems to Security Categories: Volume 1: Guide for Mapping Types of Information and Information Systems to Security Categories; Volume 2: Appendices," SP 800-60 Revision 1. http://csrc.nist.gov/publications/nistpubs/800-60-rev1/SP800-60_Vol1-Rev1.pdf and http://csrc.nist.gov/publications/nistpubs/800-60-rev1/SP800-60_Vol2-Rev1.pdf
27. NIST, "Risk Management Guide for Information Technology Systems (Sept 2012)," SP 800-30 Revision 1, http://csrc.nist.gov/publications/nistpubs/800-30-rev1/sp800_30_r1.pdf
28. NIST, "Guide for Developing Security Plans for Federal Information Systems (February 2006)," SP 800-18 Revision 1, http://csrc.nist.gov/publications/nistpubs/800-18-Rev1/sp800-18-Rev1-final.pdf
29. NIST, "Security and Privacy Controls for Federal Information Systems and Organizations," SP 800-53 Revision 4 (May 2013), http://dx.doi.org/10.6028/NIST.SP.800-53r4
30. NIST, "Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach," SP 800-37 Revision 1 (Feb 2010). http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf
31. NIST SP 800-37 and SP 800-53A (see previous notes).
32. See endnote 1.
33. For membership information, contact the Head of Membership, listed on the "Structure and governance page," www.iso.org/iso/home/about/about_governance.htm
34. All approved standards continue to be subject to periodic review to ensure timely update of industry requirements and technologies.
35. International Organization for Standardization, *Standards FAQ*. www.iso.org/iso/home/faqs/faqs_standards.htm
36. An ISMS includes any components necessary to maintain information assurance, including policies, plans, processes, roles, responsibilities, and so forth.
37. ISO/IEC 17799:2005 (superseded by ISO/IEC 27002:2005), www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=39612

38. International Organization for Standardization, "Improved ISO/IEC 17799 Makes Information Assets Even More Secure," www.iso.org/iso/home/news_index/news_archive/news.htm?refid=Ref963
39. ISO/IEC 27000:2012, http://standards.iso.org/ittf/PubliclyAvailableStandards/c056891_ISO_IEC_27000_2012(E).zip
40. Effective April 2007, ISO 17799 was superseded by ISO/IEC 27002.
41. OECD is an international cooperative committed to sharing information and developing accepted standards on economic and social issues related to trade, education, development, and science: www.oecd.org
42. IT Governance Ltd., "ISO 27001 & Information Security," www.itgovernance.co.uk/iso27001.aspx
43. IT Governance Ltd., *Information Security & ISO 27001: An Introduction*, www.itgovernance.co.uk/everything-you-want-and-need-for-iso27001.aspx
44. IT Governance Ltd., *ISO 27001 in North America*, www.itgovernance.co.uk/shop/p-1243-it-regulatory-compliance-in-north-america.aspx#.Udnt4UvDnsY
45. "Extended Mass Layoffs in 2011." Report 1039, U.S. Department of Labor and U.S. Bureau of Labor Statistics December 2012, www.bls.gov/mls/mlsreport1039.pdf (calculations based on data from Table 13, page 16).
46. C. Isidore, "Uncle Sam's outsourcing tab: $517 billion," CNNMoney June 10, 2013, http://money.cnn.com/2013/06/10/news/economy/outsourced-federal-government/index.html
47. AICPA, "SAS 70 Overview," http://sas70.com/sas70_overview.html
48. AICPA, "SSAE Overview," http://ssae16.com/SSAE16_overview.html
49. "Treadway" refers to the founding Chairman of the Commission, James C. Treadway, Jr. For a history and further details of COSO, see their "About" Web page: http://coso.org/aboutus.htm
50. AICPA, "ISAE 3402 Overview," http://isae3402.com/ISAE3402_overview.html
51. Ernst & Young, "Implementing and maintaining ISAE 3402" (2013).
52. Pub. L. 107-204, July 30, 2002, 116 Stat. 745; see *Legal Information Institute* entry for 15 USC 7201 ff, www.law.cornell.edu/topn/sarbanes-oxley_act_of_2002
53. M. Jickling, "The Enron Collapse: An Overview of Financial Issues." Congressional Research Service Report RS21135, February 4, 2002, http://fpc.state.gov/documents/organization/8038.pdf
54. M. Jickling, "Accounting Problems Reported in Major Companies Since Enron." Congressional Research Service Report RS21269, January 2002, http://congressionalresearch.com/RS21269/document.php
55. 15 USC 7241 (Section 302) (civil provision), www.law.cornell.edu/uscode/text/15/7241 and 18 USC 1350 (Section 906) (criminal provision), www.law.cornell.edu/uscode/text/18/1350
56. World Stock Exchanges, www.world-stock-exchanges.net/usa.html
57. MetricStream, "Ten Years On, Approach to SOX Shifts from Compliance to Better Governance," press release, July 24, 2012, www.metricstream.com/pressNews/pr_385_Ten-years-sox.html
58. A *significant business process* is one that accounts for a minimum portion of an organization's balance sheet. The actual threshold for inclusion will vary by organization.

59. Committee of the Sponsoring Organizations of the Treadway Commission, www.coso.org
60. Public Company Accounting Oversight Board (PCAOB). http://pcaobus.org/Pages/default.aspx
61. PCAOB Auditing Standard No. 5 superseded Auditing Standard No. 2 as of 15 November 2007. http://pcaobus.org/Standards/Auditing/Pages/Auditing_Standard_5.aspx
62. "Internal Control—Integrated Framework (2013)," COSO, May 2013. Available from "Guidance on Internal Control" page: http://coso.org/IC.htm with Executive Summary available at http://coso.org/documents/990025P_Executive_Summary_final_may20_e.pdf
63. History of ISACA, www.isaca.org/About-ISACA/History/Pages/default.aspx
64. COBIT 5, www.isaca.org/COBIT/Pages/default.aspx
65. COBIT documents in PDF are available to members and nonmembers of ISACA for free; printed copies are available at low cost (e.g., $50 for nonmembers and $35 for members). At the time of writing in July 2013, versions were available in Chinese, French, Spanish, and German, with other language translations under development. www.isaca.org/COBIT/Pages/Product-Family.aspx
66. As of the time of writing (July 2013), there were two dozen hits for the search on "PDIO" at the CISCO Website, http://tinyurl.com/l8fkvb6

CHAPTER 55

CYBER INVESTIGATION[1]

Peter Stephenson

55.1 INTRODUCTION **55·1**
- 55.1.1 Defining Cyber Investigation 55·2
- 55.1.2 Distinguishing between Cyber Forensics and Cyber Investigation 55·2
- 55.1.3 DFRWS Framework Classes 55·2

55.2 END-TO-END DIGITAL INVESTIGATION **55·9**
- 55.2.1 Collecting Evidence 55·10
- 55.2.2 Analysis of Individual Events 55·10
- 55.2.3 Preliminary Correlation 55·11
- 55.2.4 Event Normalizing 55·11
- 55.2.5 Event Deconfliction 55·11
- 55.2.6 Second-Level Correlation 55·11
- 55.2.7 Timeline Analysis 55·12
- 55.2.8 Chain of Evidence Construction 55·12
- 55.2.9 Corroboration 55·12

55.3 APPLYING THE FRAMEWORK AND EEDI **55·12**
- 55.3.1 Supporting the EEDI Process 55·12
- 55.3.2 Investigative Narrative 55·13
- 55.3.3 Intrusion Process 55·13
- 55.3.4 Describing Attacks 55·14
- 55.3.5 Strategic Campaigns 55·15

55.4 USING EEDI AND THE FRAMEWORK **55·16**

55.5 MOTIVE, MEANS, AND OPPORTUNITY: PROFILING ATTACKERS **55·17**
- 55.5.1 Motive 55·18
- 55.5.2 Means 55·20
- 55.5.3 Opportunity 55·20

55.6 SOME USEFUL TOOLS **55·21**
- 55.6.1 Link Analysis 55·22
- 55.6.2 Attack-Tree Analysis 55·24
- 55.6.3 Modeling 55·24

55.7 CONCLUDING REMARKS **55·25**

55.8 FURTHER READING **55·25**

55.9 NOTES **55·26**

55.1 INTRODUCTION. Cyber investigation (also widely known as *digital investigation*) as a discipline has changed markedly in the first two decades of the twenty-first century. In 1999, when *Investigating Computer Related Crime*[2] was published, practitioners in the field were just beginning to speculate as to how cyber investigations would be carried out. At that time, the idea of cyber investigation was almost completely congruent with the practice of computer forensics. Today (as this is being edited in July 2013), we know that such a view is too confining for investigations in the current digital environment.

Cyber investigation today is evolving into a discipline that is not only becoming commonplace in information technology but is also finding acceptance by both law enforcement and the forensic science community. The American Academy of Forensic Science, for example, now recognizes the forensic computer-related crime investigator as a legitimate discipline within its general category of membership.[3]

This chapter defines cyber investigation and examines some of its forensic components. It discusses a useful cyber investigation approach called *end-to-end digital investigation* (EEDI). Finally, the chapter explores some forensic tools and considers their usefulness both practically and theoretically.

55.1.1 Defining Cyber Investigation. Rogers, Brinson, and Robinson[4] approach the definition of cyber forensics in an ontology ("the question of how many fundamentally distinct sorts of entities compose the universe"[5]); it can also be viewed as a taxonomy ("the practice or principles of classification"[6]). Although their descriptive schema is not complete, it is informative. Most important, Rogers's approach recognizes that there is no simple definition of cyber investigation; ontology and taxonomy are therefore particularly useful in clarifying meaning and structure for such a field.

The *Rogers cyber forensic ontology* consists of five layers. Ontology generally contains classes, subclasses, slots, and instances. These elements also contain constraints and relations. Rogers's five layers can be considered to be a superclass with two classes and four nested subclasses. He does not define the other elements in his paper. The Rogers schema is easy to extend, however, to a proper ontology using a tool such as Protégé.[7] For our purposes, we will refer to his model as the *Rogers taxonomy*.

The core of the Rogers taxonomy as a definition resides in the relationships inherent in the various subclasses. Exhibit 55.1 presents the taxonomy as a hierarchy with two major classes, Profession and Technology.

Starting with *cyber forensics* at the top and with two classes represented in columns, the farther to the right in each class, the more granularity one adds to the characterization. Ontology supports understanding by providing a collection of characteristics that describe a concept; each class and nested subclass adds characteristics to that concept and contributes to a functional definition of cyber forensics. At the same time, concepts *not* included help to constrain the description and increase the precision of the ontological definition.

55.1.2 Distinguishing between Cyber Forensics and Cyber Investigation. *Cyber investigation* uses the tools of *cyber forensics* as part of investigative procedures. The Digital Forensics Research Workshop (DFRWS)[8] in 2001 developed a useful framework for digital investigation,[9] a version of which appears in Exhibit 55.2.

Although this framework appeared in 2001, it has not been improved on markedly in the intervening years.[10]

55.1.3 DFRWS Framework Classes. The DFRWS framework is a matrix. The columns are called *classes* and the cells are called *elements*. In the class descriptions that follow, the quoted italicized descriptions are taken verbatim from the original publication.[11]

The DFRWS framework classes contain key elements that are under constant review by the digital forensics community. The continuity among the classes is important; for example, the Preservation class continues as an element of the Collection, Examination, and Analysis classes. This indicates that *preservation* of evidence (as characterized by case management, imaging technologies, chain of custody, and time synchronization)

- Profession
  - Law
    - Enforcement
      - ✓ Evidence
      - ✓ Collection/Analysis
    - Courts
      - ✓ Laws
      - ✓ People
  - Academia
    - Research
      - ✓ Discipline definition
      - ✓ Problem solving
    - Education
      - ✓ Contributions
      - ✓ Professional outcome
  - Military
    - Offensive
      - ✓ Passive
      - ✓ Active
    - Defensive
      - ✓ Proactive
      - ✓ Reactive
  - Private sector
    - Industry
      - ✓ System admins
      - ✓ Legal contact
    - Consulting
      - ✓ Data recovery
      - ✓ Forensic analysis
      - ✓ Expert witness consultant
- Technology
  - Software
    - Analysis tools
      - ✓ Proprietary
      - ✓ Open-source tools
    - Operating systems
      - ✓ Proprietary OS
      - ✓ Open-source OS
      - ✓ Open source
    - File systems
      - ✓ Windows
      - ✓ UNIX/Linux
      - ✓ Mac
  - Hardware
    - Large-scale digital devices
      - ✓ Clusters
      - ✓ Grids
    - Small-scale digital devices
      - ✓ Cell phones
      - ✓ PDAa
      - ✓ SSD Operating software
    - Computers
      - ✓ Desktops
      - ✓ Laptops
      - ✓ Servers
      - ✓ Tablets
    - Storage devices
      - ✓ Thumb drive
      - ✓ Digital music player
      - ✓ External music player
      - ✓ External hard drives
    - Obscure devices
      - ✓ Gaming devices
      - ✓ Recording devices

EXHIBIT 55.1 Rogers Cyber Forensics Taxonomy

is an ongoing requirement throughout the digital investigative process. Thus, *preservation* is "a guarded principle across 'forensic' categories."[12] Traceability, likewise, is a guarded principle, but not across all forensic categories. The next topics discuss each of the DFRWS framework classes in more detail. Elements marked with an asterisk (*) in the discussions are *required* in all cyber investigations.

55.1.3.1 Identification Class. The DFRWS defines the Identification class in this way:

> Determining items, components, and data possibly associated with the allegation or incident. Perhaps employing triage techniques.[13]

The Identification class describes the method by which the investigator is notified of a possible incident. Since about 50 percent of all reported incidents have benign

Exhibit 55.2 DFRWS Digital Investigation Framework

Identification	Preservation	Collection	Examination	Analysis	Presentation
Event/Crime detection	Case management	Preservation	Preservation	Preservation	Documentation
Resolve signature	Imaging technologies	Approved methods	Traceability	Traceability	Expert testimony
Profile detection	Chain of custody	Approved software	Validation techniques	Statistical	Clarification
Anomalous detection	Time synchronization	Approved hardware	Filtering techniques	Protocols	Mission impact statement
Complaints		Legal authority	Pattern matching	Data mining	Recommended countermeasure
System monitoring		Lossless compression	Hidden data discovery	Timeline	Statistical interpretation
Audit analysis		Sampling	Hidden data extraction	Link	
		Data reduction		Spatial	
		Recovery techniques			

explanations,[14] processing evidence in this class is critical to the rest of the investigation. Likewise, as it is the first step in the EEDI process, it is the only primary evidence not corroborated directly by other primary evidence. Therefore, a more significant amount of secondary evidence is needed to validate the existence of an actual event.

The author has adopted the next definitions of the individual framework classes for the purposes of EEDI. The DFRWS has, as of this writing, not developed such definitions. Elements marked with an asterisk (*) are required elements within the DFRWS class. The elements of the Identification class are:

- **Event/crime detection.*** This element implies direct evidence of an event. An example of such direct evidence is discovery that a large number of credit card numbers have been downloaded from a server.
- **Resolve signature.** This applies to the use of some automated event detection system, such as an intrusion system or anti-virus software program. The system in use must make its determination (of the presence of an event of interest) by means of signature analysis and mapping.
- **Profile detection.** Like signature resolution, profile detection usually relies on some automated event detection system. However, in this instance, the event will be characterized through matching with a particular profile as opposed to an explicit signature. Signatures generally apply to an individual event. Events, however, may come together in an *attack scenario* or *attack profile*. Such a profile may consist of a number of events, a pattern of behavior, or pattern of specific results of an attack.
- **Anomalous detection.** Again, like the preceding two elements, this usually relies on a detection system. However, in the case of anomalous detection, the event is deduced from the detection of patterns of behavior outside of the observed norm. Anomalies can include the presence of unusual behavior but also, as in the classic Sherlock Holmes case of the dog that "did nothing in the night-time," the absence of expected behavior.[15]
- **Complaints.** This element relies on the direct reporting of a potential event by an observer. This person may observe the event directly or simply the end result of the event.
- **System monitoring.** System monitoring explicitly requires some sort of intrusion detection, anti-virus, or similar system in place. It is less specific than other elements requiring a specific action (e.g., anomaly, profile of signature detection) and may be used together with another element of this class.
- **Audit analysis.** This element refers particularly to the analysis of various audit logs produced by source, target, and intermediate devices.

55.1.3.2 Preservation Class. The Preservation class deals with those elements that relate to the management of items of evidence. The DFRWS describes this class as "a guarded principle across 'forensic' categories." The requirement for proper evidence handing is basic to the digital investigative process as it relates to legal actions.

The DFRWS defines this class as "[e]nsuring evidence integrity or state."

- **Case management.*** This element covers the management of the investigative process by investigators and digital forensic examiners. Typical in this element are investigator notes, process controls, quality controls, and procedural issues.

- **Imaging technologies.** This element is separate from the elements in the Collection class, in that it does not refer to specific hardware, software, or techniques. The imaging technologies element refers to the *technology* used for imaging computer media. For example, physical imaging or bitstream backup may be considered an appropriate imaging technology whereas a logical backup would not be. The term *imaging* as used here is rather broad. It encompasses not only the technology used to create an image of computer media but also the technology used to extract such items as logs from a device. In this case, the log might be extracted from a bitstream image or it might be read out of the device to a peripheral as a result of a keystroke command issued by the investigator.
- **Chain of custody.*** This element refers to the process of limiting access to and subsequent alteration of evidence. In most jurisdictions chain of custody rules require that the evidence custodian be able to account for all accesses or possible accesses to items of evidence within his or her care from the time it is collected until the time it is used in a legal proceeding,
- **Time synchronization.*** This element refers to the synchronizing of evidence items to a common time base. Since logs and other evidence are collected from a number of devices during the conduct of an investigation, it is clear that those devices can differ from each other in terms of time base. If all devices are in a single time-synchronized network, they will not, of course, differ. However, that rarely is the case, and some effort must be made to obtain a common time base for all devices. There are two approaches one might take. The first is to adjust all times on evidence to a common device. The second is to use a common time zone (TZ), such as Universal Time (UT) or Greenwich Mean Time (GMT), as a baseline. No evidence is modified. The investigator simply notes the variance of a particular log or other piece of digital evidence from the predetermined time standard. This also is referred as *normalizing* time stamps.

55.1.3.3 Collection Class. The Collection class is concerned with the specific methods and products used by the investigator and forensic examiner to acquire evidence in a digital environment. As has been noted, the Preservation class continues as an element of this class. With the exception of the legal authority element, the elements of this class are largely technical.

The DFRWS defines this class as "[e]xtracting or harvesting individual items or groupings."

- **Approved methods.*** This element refers to the techniques used by the forensic examiner or investigator to extract digital evidence. The concept of being *approved* refers to the general acceptance in courts of the techniques and training or certifications of the individual performing the evidence collection. The most rigorous test of methods and technologies is the *Daubert* test. (For more information on expert testimony, see Chapter 73 in this *Handbook*.) However, due largely to the immaturity of digital forensic science, most court tests have not had this level of rigor applied. For this reason, those elements in this class that relate to approval derive their authority from cases where the technique, technology, or product has been challenged in a court of the same level as the case in question and has survived the test.
- **Approved software.** This element addresses the specific software product used to collect evidence. The discussion of the approval process in the last bullet point

applies. There is an issue specifically involving software used for digital forensic data collection. In order for a software program to be considered approved, it must be identical in every way to the software that has survived either a *Daubert* hearing or a court challenge. That means that the software source code must be in every way identical in both instances of the program. Failing that, the program may need to undergo its own court testing. For the purposes of the framework and subsequent EEDI procedures, however, a program that has any differences (e.g., version level, bug fixes, source code changes, etc.) from the program tested originally is not considered to be approved software.

- **Approved hardware.** This element describes the hardware, if any, used to collect evidence. Usually this is not an issue unless the hardware is designed specifically for use in a digital forensic evidence collection environment. To a lesser extent, the caveats of sameness that apply to approved software apply to approved hardware. The hardware device used must in every way be identical to instances of the device that have survived court challenges. The approved hardware element does not apply to simple computers, disks, or other media used by the examiner to collect evidence unless the device was developed explicitly for digital forensic evidence collection and contains special unique features for use in that environment only.
- **Legal authority.*** The legal authority element is the only element of this class that is nontechnical. In most jurisdictions, some legal authority is required prior to extracting information from computer media. This authority could be a policy, a subpoena, or a search warrant, for example. Failure to comply with applicable laws may render the evidence collected useless in a court of law.
- **Lossless compression.** This element refers to the compression techniques, if any, used by backup, encryption, or digital signature software used to collect and/or preserve evidence. If the software program uses compression, it must be proven to be lossless, that is, to cause no change whatever upon the integrity of the evidence on which it is used.
- **Sampling.** If sampling techniques are used to collect evidence, it must be shown that the technique causes no degradation of the evidence collected, or, if it has, that the effects can be demonstrated clearly and unambiguously. It must also be shown that the sampling method is valid (generally accepted by the mathematical community) and that the conclusions that may be drawn from the sample are defined clearly.
- **Data reduction.** When techniques and/or programs (such as normalization) are used to reduce data that contain or may contain evidence, it must be shown that such techniques or programs produce valid, repeatable, provable results that do not affect, in any way whatever, the evidence being collected. For example, using data reduction directly on evidence would alter the evidence and would not be acceptable. However, using such methods or tools on a *copy* of the evidence would have no direct affect on the evidence. Its affect on the analysis of the evidence (the validity of conclusions, for example) is an issue for the Examination and Analysis classes.
- **Recovery techniques.** This element refers to the recovery of data that may contain evidence from a digital device. It specifically describes the methods used by the forensic examiner to extract evidence using approved hardware, software, and methods. Whereas the elements of approved hardware, software, and methods refer to the naming (or brief description of) the element and the connection

between the element and the appropriate court test by which it is approved, *recovery techniques* describes in detail the actual process used to recover the evidence. By extension, when nonforensic methods are used to collect information (e.g., traditional investigation methods such as interviewing), we consider these techniques also to be recovery techniques and we apply the same rules to them (e.g., approved methods, legal authority, etc.) as we would in a digital environment. However, we apply the rules in the context of the technique used.

55.1.3.4 Examination Class. The Examination class deals with the tools and techniques used to examine evidence. The DFRWS gives describes the Examination class in this way: "Closer scrutiny of items and their attributes (characteristics)." It is concerned with evidence discovery and extraction rather than the conclusions to be drawn from the evidence (Analysis class). Whereas the Collection class deals with gross procedures to collect data that may contain evidence (such as imaging of computer media), the Examination class is concerned with the examination of those data and the identification and extraction of possible evidence from them. Note that the Preservation class continues to be pervasive in the Examination class.

- **Traceability.*** This element is, arguably, the most important element in the EEDI process. It is the traceability and continuity of a *chain of evidence* throughout an investigation that leads to the credibility and correctness of the conclusions. According to the DFRWS, "[t]raceability (cross referencing and linking) is key as evidence unfolds."
- **Validation techniques.** This element refers to techniques used to corroborate evidence. Evidence may be corroborated in a variety of ways. Traditionally, evidence is corroborated by other, relevant evidence. However, digital evidence may stand on its own merit if its technical validity can be established. For example, a fragment of text extracted from an image of a computer disk may be shown to be a valid piece of evidence through various technical validation techniques. Its applicability or usefulness as an element of proof in an investigation may be open to interpretation, but that it is *valid* data would not be in dispute. A log, however, if extracted from a device that had been penetrated by a criminal hacker, would require additional corroboration (validation) to show that the hacker had not altered its contents.
- **Filtering techniques.** When dealing with evidence acquired from certain types of digital systems (such as intrusion detection systems), it is not uncommon to find that the gross data have been filtered for expediency by the system. Although many intrusion detection experts would agree that filtering at the source (the incoming data flow from sensors) is not as appropriate as filtering the display while preserving the original data, such source filtering does occur. This element requires that the investigator and/or forensic examiner determine and describe the filtering techniques used, if any, and apply the results of that description to the determination of the validity of the data as evidence. Another application of filtering is the extraction of potential evidence from a gross data collection[16] such as a bit-stream image of digital media. Some digital forensic tools use filters to extract data of a particular type, such as graphical images. This element requires that the filtering technique be defined clearly and understood by the investigator or forensic examiner. These tools may also use the filtering technique of matching a known hash value to digital items on a gross data collection. Items that match

the known hash are presumed to be the same as the item for which the hash value was originally generated. Again, the techniques and tools applied must be clearly understood by the investigator or the forensic examiner.

- **Pattern matching.** This element addresses methods used to identify potential events by some predetermined signature or pattern. Examples are pattern-based intrusion detection systems and signature-based virus checkers. When the pattern or signature is unclear, ambiguous, or demonstrates a large number of false positives or negatives, the evidence and conclusion following from it are open to challenge.
- **Hidden data discovery.** This element refers to the discovery of evidence that is hidden in some manner on computer media. The data may be hidden using encryption, steganography, or any other data-hiding technique. It may also include data that have been deleted but are forensically recoverable.
- **Hidden data extraction.** This element addresses the extraction of hidden evidence from a gross data collection.

55.1.3.5 Analysis Class.

This class is described by the DFRWS as "[f]usion, correlation and assimilation of material for reasoned conclusions."

The Analysis class refers to those elements that are involved in the analysis of evidence collected, identified, and extracted from a gross data collection. The validity of techniques used in analysis of potential evidence impact as directly the validity of the conclusions drawn from the evidence and the credibility of the evidence chain constructed therefrom. The Analysis class contains, and is dependent on, the Preservation class and the Traceability element of the Examination class.

The various elements of the Analysis class refer to the means by which a forensic examiner or investigator might develop a set of conclusions regarding evidence presented from the other five classes. As with all elements of the framework, a clear understanding of the applicable process is required. Wherever possible, adherence to standard tools, technologies, and techniques is critical.

The link element is the key element used to form a chain of evidence. It is related to traceability and, as such, is a required element.

55.1.3.6 Presentation Class.

DFRWS describes the Presentation class in this way: "Reporting facts in an organized, clear, concise, and objective manner."

This class refers to the tools and techniques used to present the conclusions of the investigator and the digital forensic examiner to a court of inquiry or other finder of fact. Each of these techniques has its own elements, and a discussion of expert witnessing is beyond the scope of this chapter. However, for our purposes, we will stipulate that the EEDI process emphasizes the use of timelines as an embodiment of the clarification element of this class. For more information on expert witness testimony, see Chapter 73 in this *Handbook*.

55.2 END-TO-END DIGITAL INVESTIGATION.

The end-to-end digital investigation process takes into account a structured process including the network involved, the attack computer, the victim computer, and all of the intermediate devices on the network. Structurally, it consists of nine steps:

1. Collecting evidence
2. Analysis of individual events

3. Preliminary correlation
4. Event normalizing
5. Event deconfliction
6. Second-level correlation (consider both normalized and nonnormalized events)
7. Timeline analysis
8. Chain of evidence construction
9. Corroboration (consider only nonnormalized events)

55.2.1 Collecting Evidence. A formal definition of the term *evidence collection* is:

> The use of approved tools and techniques by trained technicians to obtain digital evidence from computer devices, networks, and media. By *approved* we mean those tools and techniques generally accepted by the discipline and the courts where collected evidence will be presented.[17]

The collection of evidence in a computer security incident is time sensitive. When an event occurs (or when an expected event fails to occur), we have the first warning of a potential incident. An event may not be, by itself, particularly noteworthy. However, taken in the context of other events, it may become extremely important. From the forensic perspective, we want to consider all relevant events, whether they appear to have been tied to an incident or not. Events are the most granular elements of an incident.

We define an incident as a collection of events that lead to, or could lead to, a compromise of some sort. That compromise may include unauthorized change of control over data or systems; disclosure or modification of a system or its data; data or system destruction of the system; or unauthorized alterations in availability or utility or its data. An incident becomes a crime when a law or laws is/are violated.

Collecting evidence from all possible locations where it may reside must begin as soon as possible in the context of an incident. The methods vary according to the type of evidence (forensic, logs, indirect, traditionally developed, etc.). It is important to emphasize that EEDI is concerned not only with digital evidence. Gathering witness information should be accomplished as early in the evidence collection process as possible. Witness impressions and information play a crucial role in determining the steps the forensic examiner must take to uncover digital evidence.

Critical in this process are:

- Images of effected computers
- Logs of intermediate devices, especially those on the Internet
- Logs of effected computers
- Logs and data from intrusion detection systems, firewalls, and so on

55.2.2 Analysis of Individual Events. An alert or incident is made up of one or more individual events. These events may be duplicates reported in different logs from different devices. These events and duplications have value both as they appear and after they are normalized. This analysis step examines isolated events and assesses what value they may have to the overall investigation and how they may tie into each other.

55.2.3 Preliminary Correlation. The formal definition of the term *correlation* is:

> The comparison of evidentiary information from a variety of sources with the objective of discovering information that stands alone, in concert with other information, or corroborates or is corroborated by other evidentiary information.

Preliminary correlation examines the individual events to correlate them into a *chain of evidence*. The main purpose is to understand in broad terms what happened, what systems or devices were involved, and when the events occurred.

The term *chain of evidence* refers to the chain of events related in some consistent way that describes the incident. The relationship may be temporal or causal. Temporal chains are also called *timelines*. Essentially they say "This happened, *then* that happened." Causal relationships imply cause and effect. They say "That happened *because* this happened."

The term *chain of evidence* must not be used as a synonym for *chain of custody*, which this chapter defines in the discussion of the Preservation class of the DFRWS framework.

55.2.4 Event Normalizing. The formal definition of *normalization* is:

> The combining of evidentiary data of the same type from different sources with different vocabularies into a single, integrated terminology that can be used effectively in the correlation process.

Some events may be reported from multiple sources. During part of the analysis (timeline analysis, e.g.), these duplications must be eliminated. This process is known as normalizing. EEDI uses both normalized and nonnormalized events.

55.2.5 Event Deconfliction. The formal definition of *deconfliction* is:

> The combining of multiple reportings of the same evidentiary event by the same or different reporting sources, into a single, reported, normalized evidentiary event.

Sometimes events are reported multiple times from the same source. An example is a denial-of-service attack where multiple packets are directed against a target and each one is reported individually by a reporting resource. The EEDI process should not count each of those packets as a separate event. The process of viewing the packets as a single event instead of multiple events is called *deconfliction*.

55.2.6 Second-Level Correlation. Second-level correlation is an extension of earlier correlation efforts. However, at this point, views of various events have been refined through normalization or deconfliction. For example, during the process of deconfliction or normalization, we have simplified the collection of events to eliminate redundancies and/or ambiguities. The resulting data set now represents the event universe at it simplest (or nearly simplest). These events (some of which actually may be compound events, i.e., events composed of multiple subevents) now represent the building blocks with which we may build chains of evidence.

55.2.7 Timeline Analysis. In this step, normalized and deconflicted events are used to build a timeline using an iterative process that should be updated constantly as the investigation continues to develop new evidence. The entire process is iterative: Event analysis, correlation, deconfliction, and timeline analysis are repeated in sequence as required.

55.2.8 Chain of Evidence Construction. Once there is a preliminary timeline of events, the process of developing a coherent chain of evidence begins. Ideally, each link in the chain, supported by one or more pieces of evidence, will lead to the next link. That rarely happens in large-scale network traces, however, because there often are gaps in the evidence-gathering process due to lack of logs or other missing event data.

Such problems do not invalidate this step, however. Although we may not always be able to construct a chain of evidence directly, we can nonetheless infer missing links in the chain. Now, an inferred link is not evidence: It is more properly referred to as a *lead*. Leads can point us to valid evidence and that valid evidence can, at some point, become the evidence link. Thus again, the iterative process of refinement of evidence makes it perfectly reasonable to start with a chain that is part evidence and part leads and refine it into an acceptable chain of evidence.

A second approach to handling gaps in evidence is to corroborate the questionable link very heavily. If all corroboration points to a valid link, it may be acceptable. For most purposes, however, the former approach is best.

55.2.9 Corroboration. In this stage, we attempt to corroborate each piece of evidence and each event in our chain with other, independent evidence or events. For this process, we use the noncorrelated event data as well as any other evidence developed either digitally or traditionally. The best evidence is that which has been developed digitally and corroborated through traditional investigation or vice versa. The final evidence chain consists of primary evidence corroborated by additional secondary evidence. This chain will consist of both digital and traditional evidence. The overall process does not differ materially between an investigation and an event postmortem except in use of the outcome.

55.3 APPLYING THE FRAMEWORK AND EEDI. The important issue in application is evidence management. Both the framework and EEDI help the investigator manage evidence. They do not substitute for good investigative techniques. The incident may or may not be a criminal act; even if it is such an act, it may not be treated as one. No matter. The approach to investigating is essentially the same in either case. To understand how these investigative tools might apply, it is useful to begin with a generalized framework for understanding the execution of a cyber incident.

55.3.1 Supporting the EEDI Process. Experienced traditional investigators often resist a process-based approach to investigation. However, in numerous interviews with such investigators, the author determined that what these investigators do by habit and experience is almost exactly the same process discussed in this chapter. Thus, it is convenient to view the DFRWS framework and EEDI as tools that investigators can use to apply past physical investigation experience to the more complicated requirements of digital investigation.

55.3.2 Investigative Narrative. The investigative narrative typically consists of the investigator's notes. The EEDI process supports the construction of an investigation around an investigative framework. For the purposes of this chapter, we use the DFRWS framework shown in Exhibit 55.2. The framework includes the basic areas where investigative and forensic controls are required. Once the narrative is complete, it can be translated into a more structured evidence support and management process using the framework and EEDI.

For example, under the Collection class, we find reference to approved software, hardware, and methods. This indicates that the forensic software, hardware, and methods used by the investigator or digital forensic examiner must meet some standard of acceptance within the investigative community. That standard usually refers to court testing. Should the investigator or forensic examiner not adhere to that standard, the evidence collected will be subject to challenge. At critical points in the investigation, such as the collection of primary evidence, such a lapse could jeopardize the outcome materially.

Again, the framework does not necessarily alter the generalized investigative techniques of experienced investigators and forensic examiners. Rather, it adds a dimension of rigor and quality assurance to the digital investigative process. It also ties the functions of the forensic examiner and the investigator tightly together, ensuring that the chain of evidence is properly supported, developed, and maintained.

55.3.3 Intrusion Process. Today, there are many ways in which an attack can occur. There has been a shift since about 2000 in types of incidents from manual hacking to automated malware-generated attacks. The sources of these attacks vary widely over time, and there is no definitive pattern to how they work. Today's assaults typically are some form of blended attack. Examples are botnets, spam, Trojan horses, rootkits, worms, and other malware. (See Chapters 2, 13, 15, 16, 17, and 18 in this *Handbook* for more details of attack methods.)

A *blended attack* is one in which various intrusion techniques are combined to deliver a payload. The payload can be a malicious outcome, such as damaging the target computer, or it can be more subtle, such as a spybot that harvests personal information and calls home periodically to deliver its collected information.

Regardless of the intrusion technique—manual hacking or automated malware—there is a generalized pattern of activities associated with a cyber attack. These steps are:

- **Information gathering.** Information gathering usually does not touch the victim. During this phase, the attacker selects his or her target, researches it, and determines such things as expected range of IP addresses exposed to the Internet. There is an exception to not touching the victim: Sometimes an attacker performs a scan of a block of IP addresses to determine whether there are any vulnerable devices in the IP block. This is what security professionals usually call a *script-kiddie* attack. The usual use for this step, however, is to select a victim for a targeted attack.
- **Footprinting.** Footprinting is the act of scanning a block of addresses owned by a victim for those that are online and may be vulnerable. The attacker is obtaining the *footprint* of the victim's presence online.
- **Enumerating.** *Enumerating* refers to the analysis of exposed addresses for potential weaknesses. In this step, the attack may grab banners to determine the versions

of various services running on the target. The objective is to understand what the victim is running on exposed machines. Once that information is gathered, the attacker can select potentially successful exploits based on published vulnerabilities. Although this process will leave a footprint on the target, it is likely to remain hidden in thousands of lines of logs (as in the footprinting phase). If these two steps (footprinting and enumerating) are performed carefully, they will not trigger most intrusion detection systems.

In a fully automated attack, such as one resulting from some form of social engineering (phishing, e.g.), these two steps are replaced by research that focuses on people and their access methods (e.g., email addresses). (For more information on social engineering, see Chapter 19 in this *Handbook*; for more on phishing, see Chapter 20.)

- **Probing for weaknesses.** In this step, the attacker may perform a vulnerability scan or some more stealthy probes to determine specific exploitable vulnerabilities. In an automated attack, this phase may consist of attempting to seduce large numbers of targets to give up important information that can be used in the next step. In this case, the probes consist of social engineering attempts.
- **Penetration.** This is the entry phase. It is essentially the same no matter if the attack is manual or automated. Only the penetration technique changes. The objective always is to reach into the victim system.
- **Backdoors, Trojans, rootkits, and so on.** In this phase, the attacker deposits the payload. The objective may be to reenter the target at some later date, plant some malware, or take the system down with a denial-of-service attack. Most experts seem to agree that denial-of-service attacks are on the decline. The objective of today's intruders is more frequently stealing information than damaging systems. (For more about denial-of-service attacks, see Chapter 18 in this *Handbook*.)
- **Cleanup.** In a traditional hack, this step is where the attacker removes tools from the target, alters logs, and performs other actions to cover his or her tracks and obscure his or her presence from system administrators and forensic analysts. In an automated attack, the tasks are similar and have similar objectives, but the techniques are a bit different and the forensics required to find evidence are a bit different.

55.3.4 Describing Attacks.

Three factors lead to a successful attack:

1. The attacker can access the target.
2. There is a significant vulnerability to attack, and that vulnerability can be identified by the attacker.
3. The attacker needs to have effective command and control and the prospects for attribution are minimized.

It is important to be able to describe a cyber event in clear terms. There are several attack taxonomies and descriptive languages available to help do that. However, these tools, while useful, are inconsistent with each other (there is no universally accepted descriptive language or taxonomy) and are often unnecessarily complicated for smaller investigations. The author recommends Howard's common language[18] as a good starting point for something simple and concise. Before applying a common language, however, there are several questions that need to be asked.

- **Description of the attack.** This is a narrative describing how the attack appeared when analyzed: What is happening? What service is being targeted? Does the service have known vulnerabilities or exposures?
- **Type of attack.** Is it benign, an exploit, a denial-of-service attack or reconnaissance?
- **Attack mechanism.** This is a narrative describing how the attack was carried out: How did the attacker do it?
- **Correlations.** How does this attack compare with other similar attacks? Are there any other attacks happening at the moment that might help trace/explain the attack?
- **Evidence of active targeting.** Is the attack leveled against one or more specific targets, or is it a generic *blast*?
- **Severity.**

  Severity = (Target Criticality + Attack Lethality) − (System Countermeasures + Network Countermeasures)

  Severity formula values usually are from 1 to 5, with 5 representing the highest measure. These computations are for heuristic purposes only and do not purport to be rigorous metrics.

In answering these questions, the analyst may apply the common language (or other formal approach) as appropriate.

There are also 10 generalized questions about the incident that need to be resolved early in the investigation. These can form an informal template for early interviews.

1. What is the nature of the incident?
2. How can we be sure that there even was an incident?
3. What was the entry point into the target system? Was there only one?
4. What would evidence of an attack look like? What are we looking for?
5. What monitoring systems were in place that might have collected useful data before and during the incident?
6. What legal issues need to be addressed (policies, privacy, subpoenas, warrants, etc.)?
7. Who was in a position to cause/allow the incident to occur?
8. What security measures were in place at the time of the incident?
9. What nontechnical (business) issues may have impacted the success or failure of the attack?
10. Who knew what about the attack, and when did they know it?

55.3.5 Strategic Campaigns. The steps in an attack are, largely, tactical. This is what may be expected from an attacker during an individual attack. Often the tactical attacks are part of larger strategic campaigns. Examples are spam, identity theft carried out by organized groups, and politically motivated cyber war. There are three important differences between a tactical attack and a strategic campaign:

1. Single objective versus ongoing objectives
2. Low-hanging fruit versus sustained effort to penetrate
3. Trivial versus complicated targets and objectives

Attacks may be part of campaigns. However, just because a company experiences a tactical attack does not mean that it is part of a sustained campaign. Campaigns have distinct phases:

1. Mapping and battle space preparation
2. Offensive and defensive planning
3. Initial execution
4. Probes, skirmishes
5. Adjustment and sustainment
6. Success and termination

55.4 USING EEDI AND THE FRAMEWORK. Once we understand how attackers attack, we can see how EEDI and the framework can help us. The first, and most obvious, application of EEDI is the application of the investigative process to identify the results of each of the attack phases above. The act of collecting evidence, for example, may lead us to performing a forensic analysis of the target computer. When we do that, we will want to refer to the framework to ensure that we have not missed some important task. We may apply each step of the EEDI process to each step of the attack process to gather, correlate, and analyze evidence of an attack.

- **Collecting evidence.** Perform appropriate log collection, forensic imaging, and so on. Conduct interviews. There are 10 classes of evidence that the investigator should collect:
  1. Logs from monitoring devices
  2. Logs from hosts and servers
  3. Firewall and router (especially edge router) logs
  4. Interviews with involved personnel
  5. Interviews with business and technical managers
  6. Device configuration files
  7. Network maps
  8. Event observation timelines
  9. Notes of relevant meetings
  10. Response team notes and observations
- **Analysis of individual events.** Consider each of the events revealed by the collection and analysis of evidence in step one. Look for evidence of enumeration, penetration, and cleanup.
- **Preliminary correlation.** Tie together events from logs so that a smaller evidence set emerges. Consider each step of the attack process that applies to the device or system under analysis. Generate a straw-man chain of evidence.
- **Event normalizing.** The first and most important aspect of normalizing is time stamping. It is unlikely that all of the involved devices are on the same time tick. In this step, some allowance for differences in time zones and clocks must be made. Other normalizing tasks include combining different reports of the same events.
- **Event deconfliction.** This simply is an extension of event normalizing.

- **Second-level correlation.** This is a further refinement of the evidence and the straw-man chain of evidence. Here an important task is correlating log events, media forensic results, and results of interviews into a clear chain of evidence. This chain should be both causal and temporal, although the emphasis on the temporal chain is in the next step. This step is distinguished from the next two steps by level of detail. This correlation should offer significant detail. The next steps refine that detail making the results more comprehensible to laypeople.
- **Timeline analysis.** Refine the level of detail in the temporal chain of evidence.
- **Chain of evidence construction.** Refine the level of detail in the causal chain of evidence.
- **Corroboration.** Apply as much corroboration to evidence associated with each step of the attack process as you can.

55.5 MOTIVE, MEANS, AND OPPORTUNITY: PROFILING ATTACKERS.

Attacks, even automated ones, are performed by people. In order to understand attacks, we must understand the people who deliver them. We call those people *threat agents* because an attack simply is the delivery of a threat against a target. The act of delivering a threat is described concisely by Jones.[19] His model appears in Exhibit 55.3.

According to Jones, the threat agent first must have the capability of delivering the threat: The attacker must have access to the threat and know what to do with it. An example is an attacker who delivers a virus to a target system. The virus is the threat, and the attacker is the threat agent.

The threat agent then must be motivated to carry out the attack. However, just because motivation exists does not mean that a wily attacker will deliver the threat immediately. Usually, there needs to be a catalyst which tells the attacker that, for whatever reason, the time has come to attack.

The next step in the process, access, is the first place that we can manage the attack. Denying the threat agent access to the target denies the threat agent the ability to deliver the threat. However, assuming that the attacker can obtain access, the effectiveness of

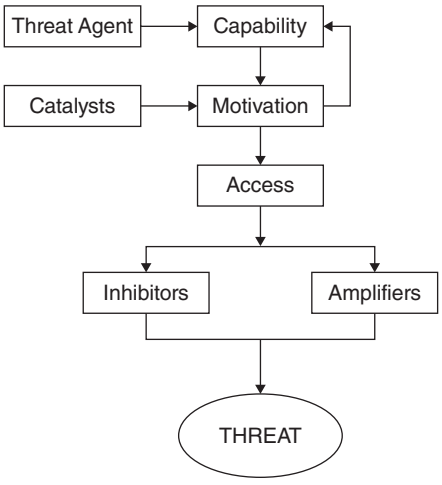


EXHIBIT 55.3 Jones's Threat Delivery Model

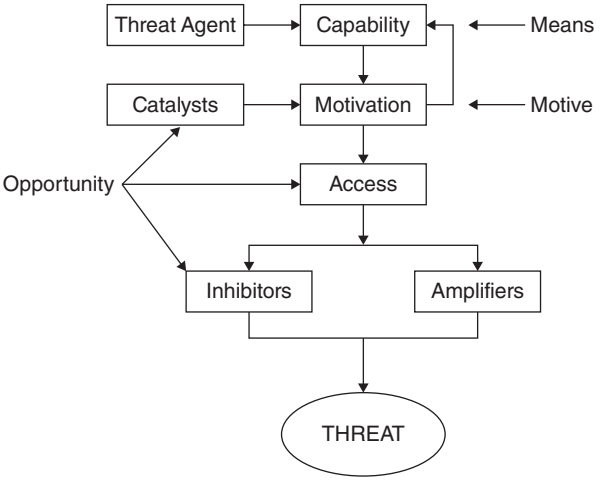


EXHIBIT 55.4 Jones's Model Showing Motive, Means, and Opportunity

the attack is now determined by the amplifiers and inhibitors present in the target system. These steps have direct analogs in the notion of motive, means, and opportunity, as Exhibit 55.4 shows.

55.5.1 Motive. Motive is critically important to understanding an attack. Why an attacker attacks often helps us understand the attack and, ultimately perhaps, who the attacker is. Traditional profiling matches evidence to known perpetrator profiles to understand the perpetrator. We do the same thing in cyber investigation.

An attacker with the motive of stealing personally identifiable information for the purpose of identity theft or credit card fraud may have an entirely different motive from a disgruntled employee who performs the same act to get even with his or her employer. In the former case, the attacker is motivated strictly by financial gain. In the latter, the attacker is motivated by a need for revenge.

Understanding the motivation helps us to determine the level of threat. In the former case, since the motivation is financial, the threat is significant and potentially far reaching because the attacker could attempt to maximize financial return. In the latter case, the threat could be somewhat less because, having stolen the information, the attacker might be likely to publicize his or her act in order to get the revenge desired. This might result in extortion attempts or leaks to the media but not in the attempt to sell or profit from the theft through sale of the information.

Additionally, analysis of motivation can help determine if the attacker is likely to be an individual or is part of a group. Groups tend to have somewhat different motivations from individuals, and they manage their attacks differently. There are several classes of threat agents, and each one has its own unique motive set. A summary appears in Exhibit 55.5.

The FBI, prior to 1995, described motive in general terms in its adversarial matrix.[20] Although this matrix needs updating for current trends and technology, it still is amazingly accurate and provides useful guidelines to the investigator. The behavioral characteristics appear in Exhibit 55.6, which is drawn from a British source.[21] This equates, approximately, to motive.

- States
  - Theft of intellectual property
  - Disruption
- Organized crime and terrorists
  - Money laundering
  - Theft of trade secrets for resale
  - Disruption
- Competitors
  - Competitive advantage
- Cyber mercenaries
  - Personal gain
- Gangs, new-age gangs
  - Financial gain
- Lone hackers and hacker "clubs"
  - Peer respect and power
  - Occasionally knowledge
- Disgruntled employee
  - Revenge
  - Personal gain

EXHIBIT 55.5 Types of Threat Agents and Their Motivations

EXHIBIT 55.6 Adversarial Matrix Behavioral Characteristics

Category of Offenders	Motivation	Personal Characteristics	Potential Weaknesses
Crackers			
Groups	Intellectual challenge; peer group fun; in support of a cause	Highly intelligent individuals, counterculture orientation	Do not consider offenses crimes, talk freely about actions
Individuals	Intellectual challenge; problem solving; power; money; in support of a cause	Moderately to highly intelligent	May keep notes and other documents of actions
Criminals			
Espionage	Money and a chance to attack the system	May be crackers operating in groups or as individuals	Become greedy for more information and then become careless
Fraud/Abuse	Money or other personal gain; power	Same personal characteristics as other fraud offenders	Become greedy and make mistakes
Vandals			
Strangers	Intellectual challenge; power; money	Same characteristics as crackers	May become too brazen and make mistakes
Users	Revenge against organization; problem solving; money	Usually have some computer expertise	May leave audit trail in computer logs

Jones also describes a motivation taxonomy:

- Political
- Secular
- Crime
- Personal gain
- Revenge
- Financial
- Knowledge or information
- Peer recognition
- Power
- Curiosity
- Religion
- Terrorism
- Competitive advantage
- Belief

55.5.2 Means. Means is a very important part of the profile. In cyber investigation, we generally equate means directly to the tools and techniques used in the attack. Indirectly we relate means to skill level of the attacker. We live in an Internet-enabled society. That state of affairs extends to computer-related crime. Relatively unskilled threat agents can deliver moderately sophisticated attacks simply because the means is readily available on the Internet to enable their actions. Analysis of evidence for indications of skill level helps us to eliminate suspects.

A word on investigative process is in order here. Although it may seem as if the right approach to an investigation is to focus on the probable perpetrator, in reality, it is far more productive, initially at least, to focus on those who absolutely could *not* be the perpetrator. This is an extension of the Sherlock Holmes' admonition: "It is an old maxim of mine that when you have excluded the impossible, whatever remains, however improbable, must be the truth."[22] The point here is that it is far more efficient to eliminate those who could not have participated in the attack. In cyber investigations, doing this becomes even more important because of the huge universe of possible participants in a network-based attack.

The FBI adversarial matrix of operational characteristics appears in Exhibit 55.7. Exhibit 55.8 extends the operational characteristics and includes specific tools, techniques, and support required. These equate, approximately, to means.

55.5.3 Opportunity. The Jones threat model describes the notion of opportunity. Considered with means, opportunity helps determine whether an attacker is a credible threat agent. Opportunity not only reflects Jones' threat model, it includes such things as knowledge of the victim system. The insider threat is an important piece of the opportunity puzzle. In this case, an insider can be the attacker him- or herself or a confederate of the threat agent. Besides insider help, the attacker may have the support of a group of some sort, such as a hacker gang, drug gang, or organized crime group.

EXHIBIT 55.7 FBI Adversarial Matrix of Operational Characteristics

Category of Offenders	Planning	Level of Expertise	Tactics/Method Used
Crackers			
Groups	May involve detailed planning	High	Each target computers via computer networks, exchange information with other crackers and groups
Individuals	Study networks before attempts are made	Medium to high, Experience gained through social networks	Use networks but more likely to use trial and error online than to do careful research and planning, use BBSs share accounts on other systems
Criminals			
Espionage	Same characteristics as crackers	High	May contract with crackers to conduct information and data collection
Fraud/Abuse	Careful planning prior to crime	Medium to high, although is typically more experienced at fraud than at computer programming	May use more traditional intrusion methods, such as wiretapping and trapdoors, will break into systems using basic methods
Vandals			
Strangers	Not much planning, more a crime of opportunity	Varies	Looks around until able to gain access to system
Users	May involve detailed planning and execution	Varies, may have high level of expertise	Trapdoors and Trojan horse programs, data modification

55.6 SOME USEFUL TOOLS. The usual tools for digital investigation include tools for:

- Computer forensic imaging and analysis
- Network forensic/log aggregation and analysis
- Malware discovery
- Media imaging (without analysis)
- Network discovery
- Remote (over-the-network) computer forensic analysis and imaging

These tools are well known in the cyber forensic community and do not warrant further elucidation here.[23] Specific products come and go so it is not appropriate to provide that level of detail here.

EXHIBIT 55.8 FBI Adversarial Matrix Resource Characteristics

Category of Offenders	Training Skills	Minimum Equipment Needed	Support Structure
Crackers			
Groups	High level of informal training	Basic computer equipment with modem	Peer group support
Individuals	Expertise gained through experience	Basic computer equipment with modem	BBS, information exchanges
Criminals			
Espionage	Various level of expertise	Basic computer equipment with modem, in some cases, uses more sophisticated devices	Support may come from sponsoring intelligence agency
Fraud/Abuse	Some programming experience	Computer with modem or access to target computer	Peer group support, possible organized crime enterprise
Vandals			
Strangers	Range from basis to highly skilled	Basic computer equipment with modem	Peer group support
Users	Some computer expertise, knowledge of programming ranges from basic to advanced	Access to targeted computer	None

In addition, there are some nonstandard tools that can make the EEDI process much easier. These include:

- Link analysis
- Attack-tree analysis
- Modeling
- Statistical analysis

Unlike the standard tools, these specialized requirements bear explanation. Again, however, specific products come and go. Where examples are provided, the specific product that generated the example will be listed along with its source at the time of writing.

55.6.1 Link Analysis. In this author's view, link analysis is the single most useful technique for cyber investigators. Link analysis allows the investigator to analyze large data sets for nonobvious relationships. Link analysis is used in the investigation of complicated criminal activities, such as fraud, drug-related activities and groups,

terrorist acts and groups, and organized crime. It is used far less commonly in analysis of cyber crime, but it should be integrated into cyber investigations.

The core theory behind link analysis is that pairings of related items (e.g., people and addresses, source and destination IP addresses, hacker aliases and real names, etc.) can be analyzed to find relationships that are not immediately obvious. We do that by testing for associations in suitable pairs of data sets.

An example of such an analysis is an analysis of hackers and groups. We might start by pairing hacker aliases and real names. These can be collected in a simple spreadsheet in columns with the first column as hacker aliases. Through research, the real names of some of the hackers may be discovered and be placed into their own column next to their associated aliases. Next, research may reveal hacks associated with individual aliases or, perhaps, real names. These go into their own column in the same row with their associated hackers. Finally, perhaps the groups with which the hackers are associated can provide useful connections. These groups are researched and placed in their own column on the rows associated with the hackers who are members.

The link analyzer accepts the input two columns at a time and iteratively analyzes the relationships. Some of the things that the analysis might reveal are multiple aliases for the same hacker, multiple hackers participating in the researched hacks, and membership in various hacker groups. In turn, membership in a particular hacker group may imply participation in individual hacks by members of the group who were not initially reported as being involved.

A more difficult problem to solve is the relationship between attack source IP addresses and target addresses. Often these relationships are *clustered*. In link analysis language, a *cluster* is

> a group of entities that are bound more tightly to one another by the links between them than they are to the entities that surround them.[24]

Since the entities within a cluster are more tightly interconnected within the cluster than they are with entities outside of the cluster, these clustered entities usually represent activities of interest.

Applying this concept to a large log from a firewall, we might take all of the denied connections, place their source and destination addresses in the link analyzer, and look for relationships. That might yield a complicated pattern such as the one shown in Exhibit 55.9.

This map can be simplified by performing a cluster analysis as shown in Exhibit 55.10.

The clustering operation on the data in Exhibit 55.9 reveals the cluster in Exhibit 55.10. This simplified analysis in turn shows two interesting relationships:

1. It quickly identifies the IP addresses that are broadcasting.
2. It clearly shows several IP addresses that are receiving a lot of traffic as well as the sources of that traffic.

The cluster map in Exhibit 55.10 is far easier to read than the map in Exhibit 55.9. Considering that the original log had several thousand entries, this is a much improved way to see nonobvious relationships; it is much easier than reading the log data or manipulating a spreadsheet.

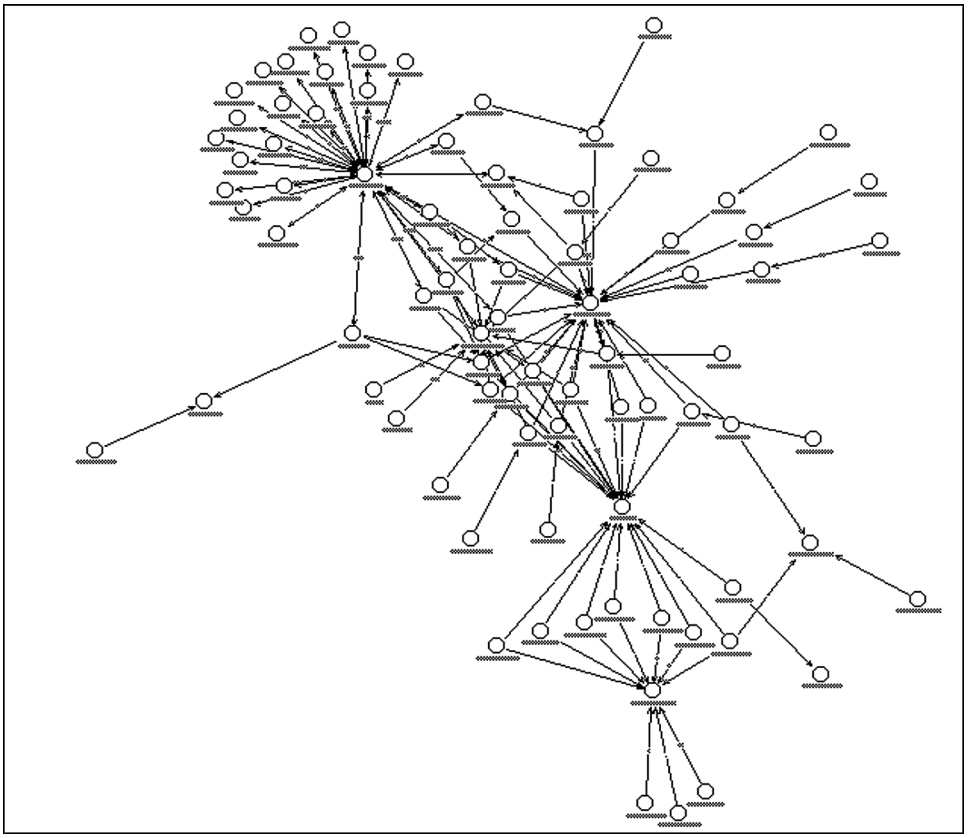


Exhibit 55.9 Link Analyzer Relationship Map of Source and Destination IP Addresses

55.6.2 Attack-Tree Analysis. *Attack-tree analysis* is a technique for analyzing possible attack scenarios on a compromised network. Its purpose is to hypothesize the attack methodology and calculate the probabilities of several possible scenarios. Once these probabilities are calculated, the investigator can hypothesize where to look for evidence. In a complicated attack on a large distributed network, having a set of these models on hand can be a big help. The analyst prepares various attack-tree models in advance and then modifies them to be consistent with the facts of a particular incident. This limits the number of high-probability attack paths and simplifies the evidence discovery process.

55.6.3 Modeling. *Attack modeling* is a broad term that refers to all of the techniques used to simulate attack behavior. The author uses Coloured Petri Nets (CPNets).[25] CPNets are a graphical representation of a mathematical formalism that allows the modeling and simulation of attack behavior on a network. To use CPNets effectively, the investigator must understand the enterprise in terms of its *security-policy domains*. The author defines security-policy domains in this way: "A security policy domain, $E_{p,}$ consists of all of the elements, **e,** of an enterprise that conform to the same security policy, **p**."[26]

The policy domains are represented along with the communications links between them in a graphical chart. Once the mathematical definitions of the domains and the

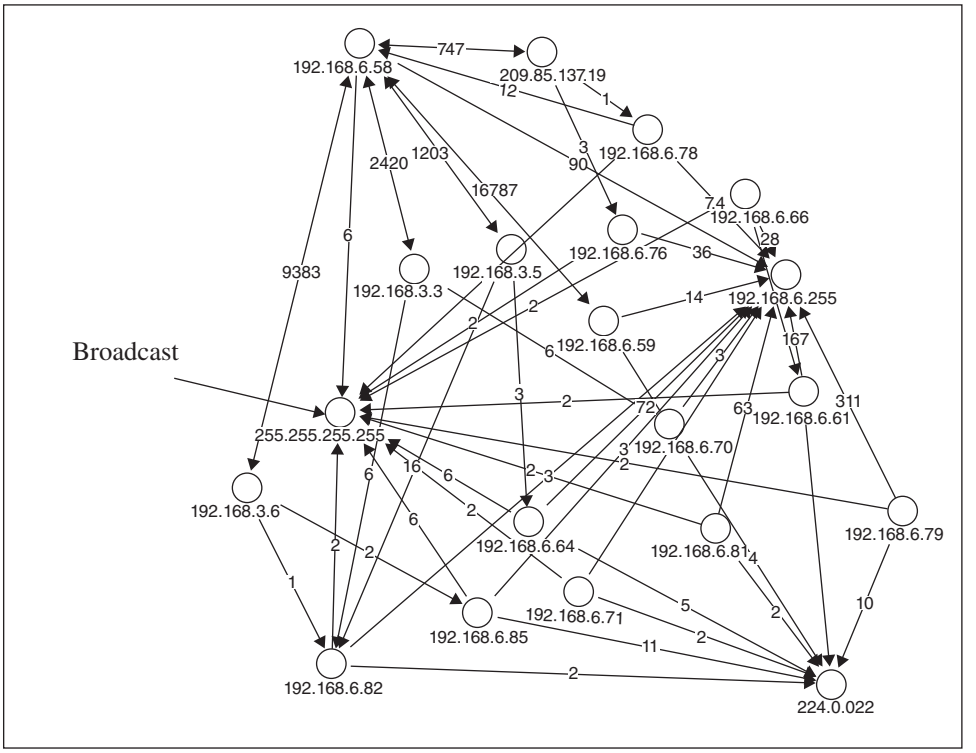


EXHIBIT 55.10 Map in Exhibit 55.9 Reduced through Cluster Analysis

communications conditions that exist between them have been defined, the behavior of the attack can be simulated. The simulations assist the investigator in tracing the attack source.

55.7 CONCLUDING REMARKS. In this chapter, we discussed a structured approach to cyber investigation along with some of the tools and techniques that support such an approach. Nordby[27] suggests that reliable methods of inquiry possess the characteristics of "integrity, competence, defensible technique and relevant experience." The nature of cyber investigation and cyber forensics supports this through rigor, underlying models and frameworks, and investigators who adhere to those characteristics.

55.8 FURTHER READING

Brown, C. L. T. *Computer Evidence: Collection & Preservation*. Hingham, MA: Charles River Media, 2005.

Carrier. B. *File System Forensic Analysis*. Upper Saddle River, NJ: Addison-Wesley, 2005.

Carvey, H. *Windows Forensics and Incident Recovery*. Upper Saddle River, NJ: Addison-Wesley, 2004.

Farmer, D., and W. Venema. *Forensic Discovery*. Upper Saddle River, NJ: Addison-Wesley, 2005.

Jones, K. J., R. Beijtlich, and C. W. Rose. *Real Digital Forensics: Computer Security and Incident Response*. Upper Saddle River, NJ: Addison-Wesley, 2005.

Marcella, A., and D. Menendez. *Cyber Forensics: A Field Manual for Collecting, Examining, and Preserving Evidence of Computer Crimes*, 2nd ed. Boca Raton, FL: Auerbach, 2007.

Steel, C. *Windows Forensics: The Field Guide for Corporate Computer Investigations*. Hoboken, NJ: John Wiley & Sons, 2006.

55.9 NOTES

1. Much of this chapter is excerpted with minor edits from P. Stephenson, "Structured Investigation of Digital Incidents in Complex Computing Environments," PhD diss., Oxford Brookes University, 2004, available in the British Library.
2. P. Stephenson and K.Gilbert, *Investigating Computer Related Crime* (CRC Press, 1999).
3. American Academy of Forensic Science, 2008: http://www.aafs.org/membership
4. M. Rogers, A. Brinson, and A. Robinson, "A Cyber Forensics Ontology: Creating a New Approach to Studying Cyber Forensics," *Journal of Digital Investigation* (2006): S37–S43.
5. *Microsoft® Encarta® 2008*, "Metaphysics" [DVD] (Microsoft Corporation, 2008).
6. *Microsoft® Encarta® 2008 Dictionary* [DVD] (Microsoft Corporation, 2008).
7. Open-source ontology editor available from http://protege.stanford.edu/
8. Digital Forensic Research Workshop: http://dfrws.org/
9. DFRWS attendees, "A Road Map for Digital Forensic Research," Report from the First Digital Forensic Research Workshop, August 7–8, 2001, Utica, NY. Air Force Research Laboratory, Rome Research Site, Rome, NY, http://dfrws.org/2001/dfrws-rm-final.pdf
10. There have also been attempts to relate cyber investigation to the generalized Zachman Enterprise Architecture Framework: J. A. Zachman, "A Framework for Information Systems Architecture," *IBM Systems Journal* 26, No. 3 (1987), http://www.zachmanframework.com/images/ZI_PIcs/ibmsj2603e.pdf
11. DFRWS attendees, "A Road Map for Digital Forensic Research."
12. DFRWS attendees, "A Road Map for Digital Forensic Research."
13. The DFRWS definitions in these sections are from Digital Forensics Research Workshop, "Day 1-DF-Science, Group A Session 1 (D1-A1), Digital Forensic Framework," Work notes for the 2003 Digital Forensics Research Workshop, August 6, 2003. For more information on the DFRWS Framework, see DFRWS attendees, "Road Map for Digital Forensic Research."
14. Author's experience over 25 years of conducting incident response.
15. A. C. Doyle, "Silver Blaze." In: *The Memoirs of Sherlock Holmes* (Newnes, 1893). http://www.gutenberg.org/files/834/834-h/834-h.htm
16. A gross data collection is a file or files containing data collected from a digital source that may contain individual evidentiary data.
17. This and other formal definitions appear in Stephenson, "Structured Investigation of Digital Incidents in Complex Computing Environments."
18. See Chapter 8, Sections 8.1 to 8.6, in this *Handbook*.
19. Andrew Jones describes a cyber threat delivery model in his paper "Identification of a Method for the Calculation of Threat in an Information Environment," unpublished (April 2002, revised 2004). http://www.researchgate.net/publication/

228957980_Identification_of_a_method_for_the_calculation_of_threat_in_an_information_environment

20. D. Icove, K. Seger, and W. VonStorch, *Computer Crime: A Crimefighter's Handbook* (O'Reilly & Associates, 1995.
21. Stephenson, "Structured Investigation of Digital Incidents in Complex Computing Environments."
22. A. C. Doyle, *The Sign of Four* (Blackett, 1890). http://www.gutenberg.org/ebooks/2097
23. For more details, see the recommendations for further reading.
24. i2 Limited "About Clusters," Analyst's Notebook v6 Help file, August 9, 2005; http://www-03.ibm.com/software/products/us/en/analysts-notebook/
25. Coloured Petri Nets at the University of Aarhus: http://www.daimi.au.dk/CPnets
26. P. Stephenson and P. Prueitt, "Towards a Theory of Cyber Attack Mechanics," First IFIP 11.9 Digital Forensics Conference, 2005. http://www.ontologystream.com/gFT/Towards a Theory of Cyber Attack Mechanics.PDF
27. S. H. James and J. J. Nordby, *Forensic Science: An Introduction to Scientific and Investigative Techniques* (CRC Press, 2003).

INTRODUCTION TO PART VI

RESPONSE AND REMEDIATION

What are the options when security breaches or accidents occur? How do we prepare for trouble so that we can minimize the consequences and respond quickly and effectively? This part includes these chapters and topics:

56. **Computer Security Incident Response Teams.** Planning and rehearsing responses to a wide variety of security problems—in advance instead of on the fly
57. **Data Backups and Archives.** The essential tool for all forms of recovery, with new material on cloud backups
58. **Business Continuity Planning.** Systematic approach to analyzing the priorities for orderly recovery when anything interrupts the smooth operation of the organization
59. **Disaster Recovery.** Planning for rapid, cost-effective return to normal after a crisis is over
60. **Insurance Relief.** Using modern insurance services to reduce the consequences of disasters
61. **Working with Law Enforcement.** Establishing relations with all levels of law enforcement before there is a crisis, and coordinating efficiently and effectively to support investigation and prosecution of criminals

CHAPTER 56

COMPUTER SECURITY INCIDENT RESPONSE TEAMS[1]

Michael Miora, M. E. Kabay, and Bernie Cowens

56.1 OVERVIEW 56·2
- 56.1.1 Description 56·3
- 56.1.2 Purpose 56·3
- 56.1.3 History and Background 56·4
- 56.1.4 Types of Teams 56·6

56.2 PLANNING THE TEAM 56·7
- 56.2.1 Mission and Charter 56·7
- 56.2.2 Establishing Policies and Procedures 56·9
- 56.2.3 Interaction with Outside Agencies and Other Resources 56·10
- 56.2.4 Establish Baselines 56·12

56.3 SELECTING AND BUILDING THE TEAM 56·12
- 56.3.1 Staffing 56·12
- 56.3.2 Involve Legal Staff 56·13

56.4 PRINCIPLES UNDERLYING EFFECTIVE RESPONSE TO COMPUTER SECURITY INCIDENTS 56·13
- 56.4.1 Baseline Assumptions 56·13
- 56.4.2 Triage 56·15
- 56.4.3 Technical Expertise 56·16
- 56.4.4 Training 56·16
- 56.4.5 Tracking Incidents 56·17
- 56.4.6 Telephone Hotline 56·21

56.5 RESPONDING TO COMPUTER EMERGENCIES 56·22
- 56.5.1 Observe and Evaluate 56·22
- 56.5.2 Begin Notification 56·23
- 56.5.3 Set Up Communications 56·23
- 56.5.4 Contain 56·23
- 56.5.5 Identify 56·23
- 56.5.6 Record 56·23
- 56.5.7 Return to Operations 56·24
- 56.5.8 Document and Review 56·24
- 56.5.9 Involving Law Enforcement 56·24
- 56.5.10 Need to Know 56·25

56.6 MANAGING THE CSIRT 56·26
- 56.6.1 Professionalism 56·26
- 56.6.2 Setting the Rules for Triage 56·27
- 56.6.3 Triage, Process, and Social Engineering 56·29
- 56.6.4 Avoiding Burnout 56·29
- 56.6.5 Many Types of Productive Work 56·30
- 56.6.6 Setting an Example 56·31
- 56.6.7 Notes on Shiftwork 56·31

56.6.8 Role of Public Affairs 56·32
56.6.9 Importance of Forensic Awareness 56·32

56.7 POSTINCIDENT ACTIVITIES 56·32
56.7.1 Postmortem 56·33
56.7.2 Continuous Process Improvement: Sharing Knowledge within the Organization 56·34
56.7.3 Sharing Knowledge with the Security Community 56·35

56.8 CONCLUDING REMARKS 56·36

56.9 FURTHER READING 56·37

56.10 NOTES 56·37

56.1 OVERVIEW. No matter how good one's security, at some point a security measure will fail. Knowing that helps organizations to plan for security in depth, so that a single point of failure does not necessarily result in catastrophe. Furthermore, instead of trying to invent a response when every second counts, it makes sense to have a competent team in place, trained, and ready to act. The value of time is not constant. Spending an hour or a day planning, so that an emergency response is shortened by a few seconds, may save a life or prevent a business disaster.

An essential element of any effective information security program today is the ability to respond to computer emergencies. Although many organizations have some form of intrusion detection in place, far too few take full advantage of the capabilities those systems offer. Fewer still consistently monitor the data available to them from automated intrusion detection systems, let alone respond to what they see.

The key is to make beneficial use of the knowledge that something has happened, that something is about to happen, or that something is perhaps amiss. Intrusion detection systems can be costly to implement and maintain. It therefore makes little business sense to go to the trouble of implementing an intrusion detection capability if there is not, at the same time, a way to make use of the data produced by these systems.

Computer emergency quick-response teams are generally called *computer security incident response teams* (CSIRTs, the abbreviation used in this chapter) or *computer incident response team*s (CIRTs). Sometimes one sees the term "computer emergency response team" (CERT), but that term and acronym are increasingly reserved for the Computer Emergency Response Team Coordination Center (CERT/CC®) at the Software Engineering Institute of Carnegie Mellon University, as explained in Section 56.1.3 of this chapter.

CSIRTs can provide organizations with a measurable return on their investment in computer security mechanisms and intrusion detection systems. Intrusion detection can indicate that something occurred; CSIRTs can do something about that occurrence. Often their value to an organization can be felt in more subtle ways as well. Many times computer emergencies and incidents cast an organization in an unfavorable light, and they can erode confidence in that organization. Efficient handling of computer emergencies can lessen the erosion of confidence, can help speed the organization's recovery, and in some cases can help restore its image. In addition, CSIRT postmortems (see Section 56.7) can provide information for process improvement (as discussed in Section 56.7.2).

When an incident occurs, the intrusion detection system makes us aware of the incident in one manner or another. We make use of this knowledge by responding to the situation appropriately. "Appropriately" can mean something different in different situations. Therefore, a well-trained, confident, authoritative CSIRT is essential.

Intrusion detection systems are not the only means by which we learn about incidents. In a sense, every component of a system and every person who interacts with the system forms a part of the overall defense and detection system. End users are often the first to notice that something is different. They may not recognize a particular difference as an incident; however, proper awareness and training will encourage them to report such situations to those who can make a determination and act on the information.

56.1.1 Description. CSIRTs are multifaceted, multitalented groups of individuals who are specially trained and equipped to respond quickly and effectively to computer emergencies.

CSIRTs come in a variety of forms and compositions. For example, while some teams are static, established groups, others are assembled dynamically to fit a specific mission or to deal with specific emergencies. Often the most effective teams are characterized as a mixture of these two approaches. These teams generally have a standing core membership made up of both technical and nontechnical members. When a situation arises that must be addressed by the CSIRT, additional members with specific skills are added to meet the requirements of handling the incident in progress. Once the incident is resolved, the team reverts to its core membership status.

56.1.2 Purpose. CSIRTs provide the first reaction to an incident. Their immediate goal should be to take control of a situation in order to contain the scope of a potential compromise, to conduct damage control, and to prevent the possible spread of a compromise to adjacent systems. Containing the scope of compromise is also synonymous with preventing or reducing loss.

Maintaining a dedicated CSIRT at the ready 24 hours a day is a costly proposition from many perspectives. Virtually any organization today, no matter its mission, will be hard pressed to justify funding such a team only to have the team stand by awaiting an emergency. The cost of maintaining a team of highly trained resources, with only emergency response roles, is a most difficult issue to overcome. Therefore, it is important to make use of team members and their skills during nonemergency periods. In many organizations, the teams have important security and awareness roles as integral parts of their charters. Carefully selected, these additional roles can benefit team readiness while at the same time providing tangible and often visible value to the organization.

For example, some CSIRTs spend their nonemergency days and nights monitoring security issues and developments for the latest trends, threats, and countermeasures. They analyze threat data and prepare reports for various levels of the organization on such topics as virus protection, password security, and emerging technology. Members of the team spend a significant amount of time and effort developing and maintaining leading-edge technical skills. They hone their response skills and procedures through continuous training. Training often is conducted in the classroom and through dry runs using a variety of response scenarios.

Often, response teams provide and maintain awareness programs for the organization. This serves several purposes. First, awareness programs benefit an organization by pointing out risks and ways to avoid them. Next, delivering awareness programs makes the team members more visible. As a result, should something unusual or out of the ordinary occur, members of the organization are not only more likely to notice it, they are also more likely to report the information, so that it winds up in the hands of the response team.

The teams respond to emergencies or incidents. An incident does not always indicate something unwanted; it also can be something that is merely unexplained, or out of

the ordinary. A response acts not only to defend, or to fight back, or to prevent further damage, but also to discover more information or to verify facts—in essence, it is part investigation and part education. To keep in step with the rapid pace of change in technology, quick response teams must be learning constantly. These teams should strive to remain abreast of each new development and technology that impacts, or has the potential to impact, the systems under their care. Therefore, responding to incidents, whether actual attacks or benign anomalies, should be seen as opportunities to sharpen the CSIRT's skills. This need is also served by the additional responsibilities the team holds in nonemergency times, typically including ongoing security research and evaluation.

If locks and other preventive measures were foolproof, intrusion detection and incident response would be unnecessary. Banks put huge vault doors, time locks, and other seemingly impenetrable defenses into their buildings. But they recognize that these measures are insufficient to prevent completely any loss of their money or other valuables. So they also install alarm systems to detect when one of the defensive barriers has been breached. But that knowledge is of little value if no one hears the alarm or, if having heard it, does not act on the information. Therefore, organizations also put into place guards, night watchers, and others, including law enforcement, to monitor systems and to respond. CSIRTs are the response part of the secure + monitor + detect + respond equation. Because connected systems are under constant passive and active attack virtually 24 hours a day, emergency response teams are a necessary part of the security equation.

56.1.3 History and Background. Help desks became popular in the 1960s for all kinds of organizations providing services and products; as computers grew in importance in the 1970s, manufacturers and users organized helpdesks specifically devoted to supporting computer users. With the increased access to the Internet in the 1980s and the arrival of the World Wide Web in the early 1990s, help desks added email, Websites and instant messaging (chat) to the range of options for customers.[2] A natural evolution of the helpdesk was to include security-incident response planning and execution; for example, the Federal Computer Incident Response Capability (FedCIRC) was proposed in April 1996 by the National Institute of Standards and Technology (NIST), the Department of Energy's Computer Incident Advisory Capability (CIAC), and the Computer Emergency Response Team Coordination Center (CERT/CC) (discussed below).[3]

The formal coordination of information from multiple CSIRTS can be traced back to one of the more notorious computer incidents of the late twentieth century. The infamous Morris Worm incident of November 2, 1988, wreaked havoc by disabling a significant portion, by some estimates as much as 10 percent, of the Internet. As organization after organization attempted to deal with the worm, it quickly became apparent that a coordinated response to such incidents would have helped to lessen the impact and speed recovery. There was no central place to report or disseminate information about the attack. Internally, few organizations were equipped with teams dedicated to responding to such attacks. As a result, they wasted time and resources duplicating efforts to identify the source of the attack, to formulate countermeasures, and finally to eradicate the worm.

At that time, the most common response was simply to disconnect from the Internet. That same response would carry unacceptable losses of revenues, confidence, and performance for today's organizations. As more commercial organizations, government agencies, and individuals became dependent on Internet-connected systems,

the criticality of the need for teams capable of responding to emergencies quickly and effectively increased.

The Morris Worm incident and several other attacks on the Internet in November and December 1988 highlighted the need for a coordinated response to widespread computer emergencies. As a result, in December 1988, under the direction of the Defense Advanced Research Projects Agency (DARPA), security experts established the Computer Emergency Response Team Coordination Center (CERT/CC®) at the Software Engineering Institute of Carnegie Mellon University in Pittsburgh, Pennsylvania. The role of the CERT/CC was to coordinate communication among organizations during computer emergencies. Their role has since expanded dramatically to include, among other things, assisting with the establishment of other CSIRT teams, acting as a clearinghouse for threat and vulnerability data, and providing training and education programs relating to security incident handling. Many entities, both public and private, have since that time adopted CSIRT's incident-handling procedures.[4]

Since its inception, CERT/CC has provided invaluable services to the world community of Internet users and especially to system and security administrators. In addition to the archives of security alerts and incident analyses available online and via free email subscriptions, CERT/CC provides free electronic textbooks of great quality. One of these is the famous *Handbook for Computer Security Incident Response Teams (CSIRTs)* edited by Moira J. West-Brown and colleagues, and which is now in its second edition.[5] We strongly recommend this work to anyone concerned with establishing and managing a CSIRT.

West-Brown et al. describe the functions of the CSIRT in this way:

> For a team to be considered a CSIRT, it must provide one or more of the incident-handling services: incident analysis, incident response on site, incident response support, or incident response coordination.

They explain in detail all aspects of these functions, and they summarize their research on the range of services that CSIRTs actually provide, whether by themselves or in cooperation with other teams in the information technology sector. They provided a "List of Common CSIRT Services" as Table 4 in their *Handbook for CSIRTs*; we have reformatted it here:

- Reactive Services
  - Alerts and warnings
  - Incident handling
    - Incident analysis
    - Incident response on site
    - Incident response support
    - Incident response coordination
  - Vulnerability handling
    - Vulnerability analysis
    - Vulnerability response
    - Vulnerability response coordination
  - Artifact handling
    - Artifact analysis

    - Artifact response
    - Artifact response coordination
- Proactive Services
    - Announcements
    - Technology watch
    - Security audits or assessments
    - Configuration and maintenance of security tools, applications, and infrastructures
    - Development of security tools
    - Intrusion detection services
    - Security-related information dissemination
- Security Quality Management Services
    - Risk analysis
    - Business continuity and disaster recovery planning
    - Security consulting
    - Awareness building
    - Education and training
    - Product evaluation or certification

The only problematic term in this list is *artifact*, which the authors define as "any file or object found on a system that might be involved in probing or attacking systems and networks or that is being used to defeat security measures. Artifacts can include but are not limited to computer viruses, Trojan horse programs, worms, exploit scripts, and toolkits" (p. 28).

The specific combination of functions that a CSIRT will provide will be a function of personnel and budgetary resources, and of the maturity of the team. It is wise to focus a completely new CSIRT on essential services such as incident handling and analysis as their first priority. With time and experience, the team can add functions, such as coordinating with other security teams, with computer and network operations in the more proactive services, and with the security quality services that will lead to long-term reduction in security incidents and to lower damages and costs from such incidents.

For more extensive discussion of the possible factors in any computer-security incident, see Chapter 8 in this *Handbook.*

56.1.4 Types of Teams. The exact composition of a CSIRT depends on factors such as the size, type, complexity, budget, and location of its sponsoring organization. Some organizations might be able to justify and support a full-time, dedicated, in-house CSIRT with the very latest technology and training. Others might improvise teams once an incident occurs, or they might even hire outside expertise to handle computer emergencies on their behalf. Still others might use a combination of these approaches by having core CSIRT staff who can be augmented by other people as needed to manage incidents.

One viable alternative to developing in-house quick response teams is to take advantage of outsourced services in this area. Outsourced incident-handling services are becoming increasingly popular, and many security companies offer them to their

customers. In some cases, this might be the most practical option for companies that lack the resources or desire to develop in-house response capabilities. Outsourcing computer emergency quick response efforts can be an effective, albeit somewhat costly, alternative to developing in-house response teams, for both short- and long-term incident handling.

However, for many organizations, establishing their own incident response capabilities can provide significant advantages. Internal teams generally know the organization and understand its goals, issues, and requirements. Outsourced responses are often mechanical and standardized. Vendors can take longer to respond since they are normally located off site and might even be a considerable distance away. Vendors may undergo frequent staff turnover, which can mean that those assigned to respond to incidents at an organization might be unfamiliar with that organization or its mission. As a result, outsiders might require precious time to "ramp up" before dealing effectively with a current incident or situation. The connected nature of today's organizations, and modern monitoring technologies, may make up for distance in some ways, but there is no substitute for an expert on site.

Some organizations are fortunate enough to have implemented formal, standing CSIRTs whose members are dedicated primarily to monitoring systems, preventing intrusions, and responding to computer emergency incidents. These teams are superior to ad hoc or outsourced teams in their ability to respond quickly with customized procedures acting from a deep and current knowledge base.

With the increased use of outsourced cloud services, it is important for in-house CSIRTs to coordinate closely with the CSIRTs at the service providers. Members of the in-house CSIRT should contact their corresponding colleagues at the service provider before there are incidents to handle; regular communication and collaboration will support smooth handling of cases involving both organizations.

56.2 PLANNING THE TEAM. Not every organization has a CSIRT already in place; not all CSIRTs are structured and managed in the most appropriate ways for an organization's specific needs. This section presents systematic approaches for rational design and implementation of a CSIRT.

Establishing a CSIRT is a complex process that must be given careful thought and must be based on comprehensive planning. Before establishing a CSIRT, the organization needs to determine exactly what it expects to accomplish. From this, the organization can decide on specific goals for the team, and perhaps most important, it can decide on policies that apply to the team. The team should be conceived and defined in terms of the organization to which it belongs. That is, the team should be tailored to achieving a specific mission. Clarity, vision, and focus are vital planning elements that ultimately will determine the success or failure of the CSIRT. Skimping during the planning stage will ensure failure. Devoting some extra effort to planning in this stage will help improve the chances of success.

56.2.1 Mission and Charter. The CSIRT should include members from every sector of the organization; key members include operations, facilities, legal staff, public relations, information technology, and at least one respected and experienced manager with a direct line to top management. The CSIRT should establish good relations with law enforcement officials and should be prepared to gather forensic evidence. The organization should have a policy in place on how to decide whether to prosecute malefactors if they can be identified. The CSIRT should be prepared to respond not only to external attacks but also to criminal activities by insiders. Proper logging at

the operating system level and from intrusion-detection systems can be useful to the CSIRT. The CSIRT plays an important role in disaster prevention, mitigation, and recovery planning.

Organizing people to respond to computer security incidents is worth the effort not only when an incident occurs but also because the analysis and interactions leading to establishment of the CSIRT bring benefits even without an emergency. A CSIRT can provide opportunities for improving institutional knowledge, contributing to continuous process improvement and offering challenging and satisfying work assignments to technical and managerial staff, thus contributing to reduced turnover. A well-trained, professional, courteous CSIRT can improve relations between the entire technical support infrastructure and the user community. The team and its members often serve as key elements of business continuity and disaster recovery teams.

A clear, written mission and a charter establishing the CSIRT are essential to its success. These documents should establish why the team exists and what the organization expects from the team at a high level. Although the current security landscape provides compelling reasons for establishing an incident response capability, identifying organization-specific goals and expectations for the team remains an essential task.

RFC 2350, "Expectations for Computer Security Incident Response," provides a solid framework that can be used and adapted in creating the charter of a new CSIRT.[6] The document includes a template and a filled-in example to serve users. The template (Appendix D) is as follows:

> This outline summarizes in point form the issues addressed in this document, and is the recommended template for a CSIRT description document. Its structure is designed to facilitate the communication of a CSIRT's policies, procedures, and other relevant information to its constituency and to outside organizations such as other CSIRTs. …

1. Document Information
- **1.1.** Date of Last Update
- **1.2.** Distribution List for Notifications
- **1.3.** Locations where this Document May Be Found

2. Contact Information
- **2.1.** Name of the Team
- **2.2.** Address
- **2.3.** Time Zone
- **2.4.** Telephone Number
- **2.5.** Facsimile Number
- **2.6.** Other Telecommunication
- **2.7.** Electronic Mail Address
- **2.8.** Public Keys and Encryption Information
- **2.9.** Team Members
- **2.10.** Other Information
- **2.11.** Points of Customer Contact

3. Charter
- **3.1.** Mission Statement
- **3.2.** Constituency
- **3.3.** Sponsorship and/or Affiliation
- **3.4.** Authority

4. Policies
  4.1. Types of Incidents and Level of Support
  4.2. Co-operation, Interaction and Disclosure of Information
  4.3. Communication and Authentication
5. Services
  5.1. Incident Response
    5.1.1. Incident Triage
    5.1.2. Incident Coordination
    5.1.3. Incident Resolution
    5.1.4. Proactive Activities
6. Incident Reporting Forms
7. Disclaimers

Effective policies are essential for any organization. The CSIRT's mission and charter should be based on organizational policies, especially information security policies. Establishing a team without having appropriate policies in place is ineffective and can put the team at odds with its own organization. Without formal policies on which to base computer emergency response activities, the team can have no legitimate basis for deciding on courses of action that support the organization. During an incident, CSIRT decisions can be unpopular. Disconnecting systems from the Internet could prevent some, or perhaps all, of the organization from carrying out its mission. Without established policies both to define and to defend those decisions, often the team is viewed as an adversary.

Without this clear definition of mission and an idea of what can be expected from the CSIRT, internal cooperation and support for the team will be difficult to obtain and even more difficult to sustain. Without internal cooperation, the team's effectiveness will be diminished, which could exacerbate the impact of an incident or prevent the team from handling an incident in a timely manner.

The overarching goal of responding to an incident should always be to prevent further damage and to restore systems and operations to normal as expeditiously as possible, consistent with organizational policies. The CSIRT members must have a clear sense of the strategic and operational priorities of the organization, just as members of business continuity and disaster recovery teams do. Without a clear idea of what the organization expects the team to accomplish, the team is likely to waste the limited time and resources it usually has available.

For more information on business continuity planning, see Chapter 58 in this *Handbook*; for discussion of disaster recovery, see Chapter 59.

56.2.2 Establishing Policies and Procedures. As the U.S. Government Defense Information Systems Agency (DISA) training course on CD-ROM about CSIRTs succinctly puts it, "policies and procedures are not merely bureaucratic red tape."[7] They are the scaffolding on which one can establish clear understanding and expectations for everyone involved in incident response. These living, evolving documents are (quoting the CD-ROM notes) tools that provide guidance on:

- Roles and responsibilities
- Priorities
- Escalation criteria
- Response provided
- Orientation

Policies are the statements of desired goals; procedures are the methods for attaining those goals. Policies tend to be global and relatively stable; procedures can and should be relatively specific, and can be adapted quickly to meet changing conditions and to integrate knowledge from experience. Policies cannot be promulgated without the approval and support of appropriate authorities in the organization, so one of the first steps is to identify those authorities. Another step is to gain their support for the policy project.

All policies and especially CSIRT policies should be framed in clear, simple language so that everyone can understand them, and they should be made available in electronic form. Hypertext can make policies more understandable by providing pop-up comments or explanations of difficult sections or technical terms. Similarly, procedures show how to implement the policies in real terms. For example, a policy might stipulate: "All relevant information about the time and details of a computer incident shall be recorded with regard for the requirements of later analysis, and for possible use in a legal proceeding." That policy might spawn a dozen procedures describing exactly how the information is to be recorded, named, stored, and maintained through a proper chain of custody. For example, one procedure might start: "Using the Incident-Report form in the CSIRT Database accessible to all CSIRT members, fill in every required field. Use the pull-down menus wherever possible in answering the questions." Again, as the DISA CD-ROM points out, these procedures should minimize ambiguity and should help members of the team to provide a consistent level of service to the organization. A glossary of local acronyms and technical terms can be helpful as part of these procedures.

Whenever policies and procedures are changed in a way that may affect users, it is important to let people know about the changes so that their expectations can be adjusted. The DISA course recommends using several channels of communications to ensure that everyone gets the message; for example, send email, use phone and phone messages, send broadcast voicemail, announce the changes at staff meetings, and use posters and Websites.

For more information on developing and promulgating security policies, see Chapters 44 and 50 in this *Handbook*.

56.2.3 Interaction with Outside Agencies and Other Resources.

No CSIRT can operate in a vacuum in an interconnected world. At some point, teams and their sponsors will require interaction with outside agencies, and even with other CSIRTs. Rather than wait until an emergency is under way, the team should establish and document contacts such as:

Internal Contacts

- **Management.** The team should establish and maintain management contacts who hold sufficient authority to make the tough business decisions that will inevitably arise during an emergency situation.
- **Systems.** The CSIRT should have a working relationship with those responsible for operating and maintaining the organization's information systems. These contacts will be necessary to allow CSIRT members appropriate system access during an emergency response situation.
- **Applications.** As with systems personnel, the CSIRT should have preestablished contact with those who manage and maintain applications. These individuals will

be able to provide CSIRT members with access to application logs, documentation, and with access to accounts during an emergency.

- **Business units.** CSIRT members should be familiar with and have established contacts with the various organizational business units they support. Having such contacts ahead of time will facilitate decision making and will avoid delays in gaining access to appropriate business personnel during an emergency.

External Contacts

- **CERT/CC.** The CERT/CC can provide valuable advice and assistance to the response team during an attack. Knowing before an emergency whom to contact, and how to reach them, will speed the process.
- **Consultants.** Often, organizations will rely on outside consultants to augment technical skills and knowledge. Therefore, it is important that the organization be able to contact these consultants during an emergency, in the event their expertise is required to respond to or to resolve a situation. Planning for such emergency contact ahead of time will avoid delays in responding to an emergency.
- **Vendors.** Responding to an emergency may require specialized information about hardware or software features and about specifications that might be available only from a vendor. Additionally, backup systems or software may need to be acquired in order to return systems to operational status during or after an incident. As mentioned in Section 56.1.4 in this chapter, coordination with CSIRTs at vendors providing services in the cloud is also essential.
- **Law enforcement.** Law enforcement agencies today frequently have specialized units capable of assisting an organization in tracing and identifying the perpetrator of an attack. However, it is important that law enforcement contacts be made in advance so that the organization can make sound decisions about whom to notify and how best to use such assistance. For more information about coordinating with law enforcement, see Chapter 61 in this *Handbook.*
- **Utilities.** Electrical power and similar infrastructure services are essential to any organization. Power outages and similar emergencies can have a devastating effect on operations. Maintaining contact with local utilities will help the organization plan for and mitigate impact from power outages. In the event of an outage, team members will know whom to contact, and they can gather information about the cause and duration of an outage more quickly. For more information on business continuity, see Chapter 58 in this *Handbook.*
- **Internet service providers.** Typically, vital connectivity is provided to the organization by one or more Internet service providers (ISPs). In many cases, ISPs are the first line of defense against some types of Internet-originated attacks, such as distributed denial of service. In addition, if an organization is attempting to track a suspected intruder, the Internet service provider will be pivotal. It is therefore vitally important that the CSIRT have pre-established contacts within the organization's ISPs in order to avoid wasting precious response time trying to get assistance during an emergency.
- **Other CSIRTs.** Other CSIRTs may well have faced the situation or emergency that another team might be facing. Other CSIRTs can provide advice and assistance, and some may even share resources and expertise to help an organization respond to an emergency. There is also an opportunity to share knowledge and conduct joint training with other teams. Information Sharing and Analysis Centers (ISACs)

in the United States are an ideal forum for such collaboration, and the National Council of ISACs provides a central source of information for locating a local ISAC.[8]

56.2.4 Establish Baselines. To be able to spot that which is out of the ordinary, the CSIRT must determine what *normal* looks like. False incidents have occurred because the observer did not have adequate knowledge to realize that the event was actually normal. Emergency response teams called into action without well-documented baselines, or detailed activity logs, must work very hard to determine whether the event is normal.

In either case, whether the triggering event turns out to be false or genuine, resources and time will have been wasted by this identification effort. A good baseline can reduce the resources expended on false positives, and can hasten the response to real events.

56.3 SELECTING AND BUILDING THE TEAM. An effective CSIRT comprises these elements, dictated by the incident at hand:

- People
- Skills
- Knowledge
- Equipment
- Access
- Authority

The makeup of the team has everything to do with how effective and responsive it will be in an emergency. Careful selection of team members at the outset will provide for an effective, cohesive group with the right skills, authority, and knowledge to deal properly with a range of known and unknown incidents.

Frequently, the first inclination is to select the most technically knowledgeable individuals available as members of the team. Although technical ability is essential to a CSIRT, this should not be the overriding characteristic. Given aptitude and motivation, appropriate technical skills can be learned. Indeed, during the course of an incident situation, adept handlers can draw on the technical expertise of people, either internal or outside, to augment their own skills and knowledge.

Maturity and the ability to work long hours under stress and intense pressure are crucial characteristics. Integrity in the response team members must be absolute, since these people will have access and authority exceeding that given them in normal operations.

Exceptionally good communications skills are required because, in an emergency, quick and accurate communications are needed. Inaccurate communications can cause the emergency to appear more serious than it is and therefore escalate a minor event into a crisis. Conversely, proper communications can galvanize others into immediate and effective activity, without creating a panic reaction.

56.3.1 Staffing. The CSIRT may be a permanent, full-time assignment for a fixed group of experts, or it may be a part time role formed dynamically as conditions require. In either case, or for any of the intermediate arrangements, certain fundamentals will dictate the choice of staff members.

The DISA course on CSIRT management also addresses the question of the technical level required by CSIRT staff. The course authors suggest using a scale from 1 to 10, with 1 representing the novice or support staff and 10 representing the technical wizard. Individuals in the 1-to-3 technical range should be sufficient to handle the initial triage process, which involves separating service request into categories, and directing them to the appropriate team member.

Information requests can be handled by team members in the 1-to-5 range. For example, a support staff person can send out publications, while someone with greater expertise would be required to address questions about identifying spoofed email.

Team members in the 5-to-8 technical range are necessary to respond to actual incidents. This response can involve technical analysis and communicating with compromised sites, with law enforcement technical staff, and with other CSIRTs. In handling incidents that represent new attack types, it may be necessary to call "wizards" to help understand and analyze the activity.

Vulnerability handling requires the most proficient personnel, falling into the 8-to-10 range. These individuals must be able to work with software vendors, CSIRTs, and other experts to identify and resolve vulnerabilities. Many CSIRTs do not have access to this level of technical expertise.

CSIRT staff with the psychological flexibility to allow them to adapt quickly to changing requirements will do better than people who resist change or resent ambiguity. Ideally, the team will include problem solvers with an intuitive grasp of the differences between observation and assumption, hypothesis, and deduction. As always, team players committed to getting the problem solved will contribute more than people interested in acquiring personal credit for achievements. Having at least one person on the team with a penchant for meticulous note taking is a real benefit (see Sections 56.4.5.2 and 56.5.6).

56.3.2 Involve Legal Staff. As with any crisis event, every action carries with it a potential legal implication. This is especially true in a situation where an evidentiary chain may be required. Even if evidence is not a primary concern, due diligence requires that accurate records be kept of the incident costs, including response team costs, and of the scope of compromise and effect.

The corporate legal staff must play an important role in developing response team procedures, in training the response team, and in crisis resolution.

56.4 PRINCIPLES UNDERLYING EFFECTIVE RESPONSE TO COMPUTER SECURITY INCIDENTS. Some general considerations underlie effectiveness for all incident-handling teams.

56.4.1 Baseline Assumptions. The primary consideration in responding to any emergency situation must be given to preventing loss of human life. Following that, a comprehensive information security program will have included the identification and classification of the most sensitive data and systems. This classification should provide a clear prioritization of what should be protected first, in the event of an emergency. For example, a business's survival might depend on the confidentiality and integrity of some intellectual property, such as engineering diagrams. After the safety of personnel, those drawings, and the systems on which they are stored, would be the obvious first priority for protection. A network intrusion that threatened an email server located on a separate network segment would likely not warrant immediately disconnecting the entire network from the Internet. If the particular server containing the company's

engineering drawings, or the network segment on which it sits, were under direct attack, however, an appropriate first response might well be to disconnect from the Internet.

In any case, there must be an unambiguous sense among the CSIRT that those responsible for taking actions in good faith will not suffer reprisals as a result of taking those actions. For example, based on facts in evidence at one time or another, a member of the team might decide to disconnect an operational system from the Internet because it appears to be under attack or compromised. Should this turn out to be a false alarm of some sort, the individual authorizing the action should suffer no reprisal or sanction by taking what he or she believed was a legitimate action to stop, or to respond to, an attack.

With appropriate plans in place before incidents happen, recovery can be effected much more quickly and with less residual damage.

Appropriate responses depend on the systems involved, and should be documented and agreed to before those systems are connected to the Internet. If a router, or a firewall protecting the outer perimeter of the network, becomes compromised, it may be necessary to disconnect the entire system in order to contain the situation.

If a Web server in an isolated part of the network becomes compromised, though, disconnecting only that system should be sufficient. However, any business-related activities carried out using that server might no longer be available. A loss of revenue or image might result from disconnecting the server. Business units and others that depend on these systems must be made aware of, and must agree to accept, the impact of proposed responses the CSIRT might take during an incident.

Planned responses, combined with the authority and confidence to execute them, can save the organization both time and money. As an example, if every incident requires the presence of senior business managers or executives, those leaders must be taken away from their normal duties during rehearsals and during actual events. If the CSIRT leader on duty or on call finds the incident to be routine, and if the incident has been well planned for, a simple notification of the facts can be sent to appropriate senior-level personnel, leaving them free to attend to their normal duties and to participate only in a major event at their appropriate level of management.

Planned, preapproved responses can speed reaction times, enhance security, and lessen impact of a given breach or incident. In most cases, the CSIRT leader will follow a series of commonsense steps to handle an incident from identification through resolution. As the leader progresses through each step, he or she may choose from one of the preapproved responses to handle the incident, or the situation may require the involvement of other people and resources for resolution. In either case, the basic flow of events should be similar to this list:

- **Triage.** Deciding how to direct calls for help or reports of a computer security incident
- **Technical expertise.** Assembling the different kinds of knowledge required to support an effective response
- **Tracking incidents.** Ensuring appropriate documentation to save time and reduce errors
- **Critical information.** Laying the ground rules for collecting the kinds of data needed for effective decisions
- **Telephone hotline.** Establishing protocol for real-time notification and response

56.4.2 Triage. The word *triage* itself comes from a French root meaning *to sort*. In medicine, triage is "prioritization of patients for medical treatment: the process of prioritizing sick or injured people for treatment according to the seriousness of the condition or injury."[9] Similarly, anyone receiving calls about computer security incidents must be able to classify the call right away, so that the right resources can be called into play. As the DISA course on CSIRT management suggests:

> The triage process recognizes and separates
>
> - new incidents,
> - new information for ongoing incidents,
> - vulnerability reports,
> - information requests, [and]
> - other service requests.[10]

We have altered the order of the original list to reflect a decreasing rank of importance for these factors in communicating and acting on calls.

Triage is common to ordinary help desks as well as to emergency hotlines. In general, there are two models for staffing the phones for such front-line functions: the "dispatch" model and the "resolve" model.[11]

1. The *dispatcher* has just enough technical knowledge to collect appropriate information about an incident and to assign a team member for investigation; the alternative is to assign someone with more expertise to answer the phone so that response can be even faster.
2. However, the *resolve* model risks wasting resources because the more experienced staff member may end up doing largely clerical work instead of focusing on applying expertise to problem analysis and resolution.

To support triage, staff members need explicit training on data collection and priorities. They need to record: who is calling; how to reach that person; what the caller thinks is happening; what the caller has observed; how serious the consequences are; how many people or systems are affected; whether the incident is in progress or is over, as far as they know; and how the caller and others are responding. The CSIRT procedures should include guidance on assigning priorities to incidents; factors can include:

- Security classifications (e.g., SECRET or COMPANY CONFIDENTIAL data under attack)
- Type of problem (e.g., breach of confidentiality, data corruption, loss of control, loss of authenticity, degradation of availability or utility)
- Possible direct costs (e.g., personnel downtime, costs of recovery, or loss of business)
- Possible indirect costs (e.g., damage to business reputation or legal liability and so on), as appropriate for each organization

Readers may find the work of John Howard relevant for such analysis; Dr. Howard established a useful taxonomy for discussing computer security incidents that can serve as a framework for establishing priorities.[12] See Chapter 8 in this *Handbook.*

We recommend an automated system for capturing information on all calls to the CSIRT. Using keywords "helpdesk software" and also "help desk software" in an online search brings up dozens of options for such programs. Readers with modest skills in database design can also create their own using a program such as Microsoft Access, but it may take a good deal of time to improve the home-grown system so that it matches commercial or freeware versions based on extensive experience of the user community. However, with appropriate locking strategies to permit safe concurrent access and with well-designed automated reports, your CSIRT can know and control the priorities of all the open incidents under investigation at any time.

For more about management perspectives on triage, see Section 56.6.3.

56.4.3 Technical Expertise. The DISA CD-ROM course starts by classifying technical expertise in approximate ranges:

- **Low**, suitable for the triage function which involves determining who should best handle a specific call
- **Medium**, appropriate for answering requests for information
- **High**, suitable for technical problem solving
- **Expert**, suitable for handling problems that others have been unable to resolve and especially for issues involving vulnerability analysis and real-time response to attacks

As the DISA writers point out, "Vulnerability handling requires your most proficient personnel. … These individuals must be able to work with software vendors, CIRTs, and other experts to identify and resolve vulnerabilities. Many CIRTs don't have access to this level of technical expertise."

56.4.4 Training. Teams that have no experience responding to incidents are of little value to an organization. Predictably, computer emergency responses by untrained or inexperienced teams result in loss or destruction of evidence, legal exposure by failing to properly protect individual rights, and failure to properly document and learn from the experience. To be most effective, training must be iterative (learn, exercise, review, analyze, repeat) and should involve as many realistic scenarios as possible, so that the CSIRT becomes exposed to a wide variety of potential emergency situations.

56.4.4.1 Rehearse Often. Experience can be gained only by responding to incidents or through training with simulated attacks. While the time and resources required to practice responding to incidents might be costly, more costly still is the potential damage resulting from an uncontained or poorly handled breach of the system.

An excellent opportunity to practice response procedures, and to develop response teams, occurs during periodic security assessments, including penetration testing. Penetration tests simulate external and internal attacks on a system and offer a real yet controlled environment in which to exercise, train, and evaluate a CSIRT. The simulation is especially effective when an outside, independent team is engaged to conduct the penetration test and security assessment. With proper coordination, the test can provide an opportunity for the team to observe, and react to, many different types of incidents.

56.4.4.2 Perform Training Reviews. At the conclusion of any training exercise, it is important to reassemble the team as soon as possible, not only to review management's view of its performance but to reveal its own perspectives as well. Each of the participants should be asked what went right and what went wrong. Were necessary resources (information, decision makers, tools, software, equipment) unavailable when the team needed them? Did the team have, or was it able to obtain in a timely manner, the physical and system access it needed? Did it have the right documentation and access to other company personnel? Were systems for communicating among the team and with other company or external personnel adequate and efficient?

Video cameras are a useful tool for recording events during the training sessions; many employees will have such equipment available and can make inexpensive recordings that can be analyzed during the training reviews.

See Section 56.7.1 on postmortem analysis of real incidents for more ideas that can be applied to training reviews.

56.4.5 Tracking Incidents. This section focuses on some of the advantages, requirements, and tools for incident tracking. First, we establish why documentation in general is so important.

56.4.5.1 Will This Have to Be Done Again? When one of the authors (MK) joined Hewlett-Packard (Canada) Ltd. in 1980, he arrived on the job armed with a small, green, hard-covered book prominently labeled LOGBOOK in big black letters. From his first day as a member of the systems engineering organization, he wrote down what he learned; he logged how he spent his time. When he met clients, he took notes. When he installed new versions of the MPE operating system, he kept a chronological record of everything he did—including mistakes. When he taught courses, he kept a list of questions he could not answer right away.

Pretty soon, people began asking him what he thought he was doing, writing a novel?

His colleagues may have been puzzled by what they perceived as a mania for record keeping, but he was equally astonished that record keeping was not a normal part of their way of doing work. The reason for automatically kept records was his years in scientific research, where logbooks with hard covers, numbered pages, and even waterproof paper were just usual parts of doing serious work. The idea of doing anything of importance without keeping a concurrent record simply did not occur to anyone in research. One could not reproduce an experiment without knowing exactly what sequence one had used in accomplishing the steps. Even adding salts to solutions had to be done in a particular order.

So he just kept on keeping his little green logbooks.

56.4.5.2 Why Document? Documentation, far from being a sterile exercise done to conform to arbitrary requirements of nameless, faceless superiors, should be a vital part of any intellectual exercise. Documentation is simply writing down what we learn: the crucial step in human history that changed traditional cultures into civilizations. By keeping a record independent of any specific individual, we liberate our colleagues and our successors from dependence on our physical availability. Documentation is our assurance that work will continue without us; a kind of immortality, if you will.

We document what we do as a part of systematic problem solving. Writing forces us to identify the problem in words, instead of being content to define it in vague, unclear

ideas. Writing down each idea we are in the process of testing helps us notice the ideas we missed the first time we tackled the problem. Keeping notes helps us pay attention to what we are doing.

Documenting what we do also helps us during training—both our own and that of the people we are helping to learn technical skills. Trainees can review their own notes on how to do something instead of relying entirely on someone else's description. If taking notes is viewed as a chance to engage one's mind more thoroughly in what we are learning, it can be stimulating and even fun.

Finally, accurate records can be a boon in legal wrangles. In a case one of us (MK) experienced, upper management seriously considered legal procedures against a supplier for supposed breach of contract. Careful records of exactly when meetings were held, and with whom, permitted us to analyze the problem and to resolve the issues by collaboration instead of by confrontation. Such records, if kept consistently, in good times and bad, can be accepted in a court of law as evidence—but only if everything points to a steady pattern of record keeping as events unfold. Records made long after a problem occurs are worthless.

56.4.5.3 Keep Electronic Records. The best way of keeping records on specific problems is an easy-to-use database. Such records help team members remember and share information that can help in solving new problems as they arise. With easily accessible records, it is possible to solve problems without the presence of specific team members. Such shared knowledge speeds problem resolution, improves the competence of all team members with access to the knowledge base, and provides a sound basis for training and integration of new team members. Sharing knowledge can be a liberation for key members of any organization by sparing them from the sense of obligation to be present at all times; it also supports management policies that enforce security principles by requiring employees to take vacation time. For a discussion of the dangers of allowing any employee to become indispensable, see Chapter 45 in this *Handbook* on employment practices and policies.

56.4.5.4 Advantages for Technical Support and CSIRTs. Keeping track of all of technical support calls is essential for effective incident handling. Having details available to all members of the CSIRT in real time, and for research and analysis later, serves many functions:

- **Communication among team members.** Having the details written down in one place means that team members can pass a case from one to another and share data efficiently.
- **Better client service.** Callers become frustrated when they have to repeat the same information to several people in a row; a good incident-tracking system reduces that kind of irritation.
- **Documentation for effective problem solving.** A good base of documented experience can help find the right procedure and the right solution quickly.
- **Institutional memory.** When experience is written down and accessible, the organization's capacity to respond quickly and correctly to incidents improves over time.
- **Follow-up with clients.** Managers can use the incident database to prepare management reports and to follow-up with specific clients to understand and resolve difficulties or complaints.

- **Forensic evidence.** Detailed, accurate, and correctly time-stamped notes can be a deciding element in successful prosecution of malefactors.

56.4.5.5 Requirements. Some of the more obvious requirements of any incident-handling system are:

- A unique identifier for each case
- Dates and times for all events
- Who currently controls the case: It should be instantly obvious who is in charge of solving the problem
- Keywords
- Contact information: Every person in the case should be listed with office phone, cell phone, email, and fax numbers
- Handover of control: Whenever someone takes over control of the case, that handover should be noted in the record
- Technical details, including:
  - Diagnostics
  - Tests of hypotheses
- Resolution: What was the outcome? When was the case closed?
- Search facilities: Full-text search capabilities
- Knowledge base: Ability to integrate vendor-supplied entries to speed research
- Industry-standard database engine: Easy to learn, maintain, and improve
- Accept input from comma-separated value (CSV) files: Import data from other systems

56.4.5.6 Tools. There is a wide range of software available for tracking incidents. One can build one's own, but then proper documentation and training materials must also be created, because turnover is a constant problem for CSIRTs. In addition, unless analysts have experience with the CSIRT function, they are likely to miss useful features that have accumulated over the years in products used by thousands of people.[13] Well-respected open-source tools are listed in the "Further Reading" section of this chapter.

All such tools can be complex. To prevent people from fumbling about in an emergency, a budget must be established adequate for staff training in implementation of the selected tool.

56.4.5.7 Get the Global Picture. When gathering information about an incident, staff members should establish a clear picture of what people were doing when they realized that there was a problem. For example, it may be important to know that someone was accessing a rarely used account and noticed that a file was not available because someone else had it open. Those details will help to characterize the attack and to provide clues that may lead to additional valuable data. However, the CSIRT investigator should also ask why the contact was accessing the rarely used account; it takes only a minute, but getting a wider picture may give the analyst another perspective that can also lead to new clues. In the scenario just sketched, one could imagine that a system administrator had become curious about some unexpected resource utilization in a supposedly dormant account. This simple fact might lead to additional exploration of

system log files and questions about whether any other dormant accounts had sparked curiosity. So, in general, it is worthwhile to explore the situation more broadly at first, rather than driving down the very first avenue that presents itself in the initial questions.

For more information on using log files for analysis of problems, see Chapter 53 in this *Handbook*.

56.4.5.8 Distinguish Observation from Assumption. As the CSIRT member listens to the observations of other staff members, it is critically important to distinguish facts—that is, personal observations—from assumptions. Assumptions are ideas taken for granted or statements that are accepted without proof. For example, imagine the serious consequences of hearing someone say, "And so then they must have exploited a flaw in the firewall and then they…" and simply writing that assumption down as if it were a tested and validated explanation of the events. Such an assumption could profoundly distort the investigation, putting people's efforts onto the wrong track and diverting their attention from a more fruitful line of inquiry. Hearing such a statement, one should write down, "And so perhaps they exploited a flaw in the firewall and…" or "Bob thinks that they exploited a flaw in the firewall and…"

56.4.5.9 Distinguish Observation from Hearsay. Everyone has played the child's game of whispering a sentence to another person and then hearing the distorted version that come out the other end of a long chain of transmission without error correction. CSIRT staff must always distinguish between first-person observations ("I read the log file and found…") and hearsay ("Shalama read the log file and she found…"). Team members should not trust hearsay: They must check it out themselves by tracking down the source of the information. Even when someone is reporting a personal observation to a CSIRT member, it is important to weigh the cost of verifying the observation (when possible) against the consequences of branching off into the wrong part of the solution space.

56.4.5.10 Distinguish Observation from Hypothesis. Sometimes when people are careless or untrained, they do not distinguish between what they *saw* and an idea that might *explain* what they saw. In the previous example about a supposed flaw in a firewall, the person speaking seemed to take the flaw for granted; that was an assumption. A similar problem can occur when someone thinks that *maybe* there is a flaw in the firewall and then proceeds as if that were true without testing the hypothesis. "And so maybe they exploited a flaw in the firewall, so we should patch all the holes right away." Putting aside for the moment the advisability of patching holes in firewalls, merely hypothesizing an exploit does not make it true. Maybe it is a good thing to patch the firewall, but it does not follow that it is the top priority right now simply from having thought of the idea. CSIRT staff should be careful to think about what they are hearing and should note explicitly when people are proposing explanations rather than reporting facts.

56.4.5.11 Challenge Hypotheses. When CSIRT members develop hypotheses about what is happening in a breach of security, they have to do two things: see if the ideas are consistent with observation but also test those ideas to see if they are flawed.

Trying to show that an idea is *correct* is a natural response when solving problems. Especially when the clock is ticking and a critical process is stopped, the immediate

need is to get the system running immediately, no matter what it takes. Unfortunately, doing something and having the system work afterward does not automatically mean that the solution being proposed actually fixed the problem—that is the fallacy known as *post hoc, propter hoc* (after the fact, because of the fact). It is possible that what we think fixed the problem simply preceded a change of state related to some other factor. The supposed fix may have nothing to do with the solution or may be only part of the solution. Even under pressure, technical support teams with experience go beyond the immediate fix to see if there are other factors that need to be controlled for long-term stability of the system.

In analyzing the behavior of a compromised system, the CSIRT is usually less concerned with restart than with forensic analysis. Who did what to which parts of the system? What do the log files tell us about the incident? How could the attacker have gotten in? What might (s)he have changed?

For that kind of analysis, it is especially important to find ways of testing our ideas before we go down a long chain of reasoning that may be flawed at its very start. Thus, just as in quality assurance, try to come up with ways of showing that our explanation is *wrong*. If we fail to disprove a hypothesis using genuine, thoughtful, intelligent tests of our ideas, maybe we have got something useful after all. In practice, the principle teaches us to go a step further when solving problems. Instead of stopping the testing as soon as we find supportive evidence, we can make it a habit to ask "Yes, but what if… ?"[14]

56.4.6 Telephone Hotline. Users should be trained and encouraged to call the telephone hotline—usually the helpdesk line—to report anomalies or suspected breaches of security. The helpdesk operator can route the call to the appropriate persons, including the CSIRT monitor.

Returning for a moment to the advice on staffing the CSIRT as discussed in Section 56.3.1, there are some additional requirements for the people involved in the CSIRT concerning their interpersonal relations. Not only should managers look for, and ensure, adequate technical knowledge, they should also enhance interpersonal skills and disciplined work habits.

CSIRT members inevitably work with some users who are stressed by the problems they are facing. It is no help to have a technical wizard who so offends the users that they stop cooperating with the problem-resolution team. Sometimes CSIRT staff members forget that their job includes not only resolving a technical issue but also keeping the clients as happy as possible under the circumstances—and the use of the word "clients" is deliberate here.

Here are some of the most irritating responses to users we have run across in several decades of technical support and client support followed by our comments in brackets:

- "No one has ever complained about this before." [So what? If the problem is real, we should thank the user for reporting it, not make veiled criticisms that imply that the problem cannot be real.]
- "I don't have time for this now." [That is a time management problem for the CSIRT, not for the client. Take responsibility for getting the right person to take charge of the problem in real time.]
- "Why don't you try calling… ?" [Same comment as last one.]
- "That's not my problem." [Just plain rude as well as irresponsible.]

- "Why don't you reload the operating system and call me back if it happens again?" [Significant risk and time cost for the client; often the first-line suggestion of the terminally incompetent technician.]
- "Just format your hard disk and see if it happens again." [Even worse than the previous suggestion if it is just a casual suggestion to get the client off the phone for now.]
- "Don't get mad at me—I just work here." [A professional will understand that there is a difference between criticism directed at the organization or its procedures versus a direct *ad hominem* attack. The former should be taken seriously and passed on to people who can evaluate the seriousness of the criticism; the latter can be unacceptable and should be passed on to a manager who can explain to the user the need for and value of civility even under stress.]

56.5 RESPONDING TO COMPUTER EMERGENCIES. This section offers specific recommendations and comments on a structured response plan to help develop a systematic approach to CSIRTs.

56.5.1 Observe and Evaluate. A response team leader must assess the situation as quickly as possible based on available information. The leader should make a preliminary estimate of the type of incident, its scope, the people involved, and the data or systems affected, and then begin formulating first responses. This is the point at which the team leader or other responsible person orders a move from a state of standby monitoring to one of active monitoring, focused on the particular event or events. It is important to maintain standby and baseline monitoring activities during an actual incident, because the obvious event might well be a ruse designed to divert attention from a more serious attack.

If proper planning has taken place, the team leader usually will be able to direct a specific course of action in response to a particular incident. The leader can choose from a menu of planned responses while drawing on only those resources necessary to execute that particular response. Doing this minimizes the impact on staff at all levels and allows the incident to be dealt with efficiently and effectively. However, the more unique or complex the situation, the more likely it is that a complete team response may be required.

Responses, the players involved, and the audience are obviously different when considering a data center–type situation as opposed to a user-reported situation. Often, formal handling procedures are pre-established for data centers, which generally are staffed by more technical personnel. The CSIRT can expect a higher level of response from data center personnel and will likely be able to communicate instructions more concisely and with more assured compliance.

Dealing with individual users, however, requires a greater degree of sensitivity and understanding. In most cases, the CSIRT will be moving quickly and enthusiastically when handling an incident, since this is what the team has trained so long and hard to do. Individual users often are stressed or bewildered when confronted by a computer emergency incident serious enough to warrant a response team. In these situations, users tend to be nervous rather than excited or confident. This can cause communication problems, especially when a team member converses with users over a phone. Instructions become garbled, or may not be carried out exactly as desired.

Team members must be trained to communicate clearly and calmly when dealing with individual users who may not have exceptionally well-developed technical skills.

This is especially important when giving end users instructions over the phone or via some remote means. A calm, careful conversation will lessen the amount of stress the individual feels while helping to ensure that the CSIRT member's instructions are carried out properly. Proper compliance with instructions can make or break an incident investigation, especially when forensic issues are to be considered, as in the case of known or suspected criminal activity. Failure to maintain the state of an attacked system properly can thwart any subsequent attempt at a successful prosecution.

56.5.2 Begin Notification. Once the team leader establishes that an incident is in fact in progress, notification must begin to appropriate individuals within the organization, consistent with the type of situation. Notification and actions should be carried out, whenever possible, according to existing plans.

In some cases, the CSIRT leader might be able to identify the incident as one calling for a prearranged response. The leader, having the authority and confidence to carry out such a preapproved response, will notify those appropriate to the incident and carry out the contemplated actions. In other cases, the situation might not be so clear, and the notification process might include additional personnel with authority to decide on various courses of action.

56.5.3 Set Up Communications. Team members, especially when dealing with remote or multiple sites, must be able to communicate easily and securely with one another as well as with management representatives. Team members need to be able to communicate data, status updates, actions, responses, and similar events. Communications should flow securely to the designated CSIRT leader for coordination. The team leader must be able to direct and advise other team members, but the potentially sensitive nature of an incident may require that these communications be handled out of band and through secure means. "Out of band" in this case refers to communication methods that are neither part of nor connected to the system believed to be under attack. For example, in communications regarding an attack, the use of unencrypted email that might be intercepted by an attacker, or by other unauthorized parties, should be avoided whenever possible.

56.5.4 Contain. The CSIRT's next course of action is to contain the incident. The goal is to limit the scope of any compromise as much as possible, by isolating the system under attack from other systems in order to prevent the problem, attack, or intrusion from spreading. Containment might involve steps such as disconnecting systems from the Internet. However, doing so might limit the organization's ability to catch an intruder who is currently active on the system. The priority level assigned to intruder identification and prosecution is a part of the mission and charter of the team, modified by the specific action plans in use for a particular incident.

56.5.5 Identify. Once the team has taken steps to contain the incident as much as possible, it should focus on identifying exactly what happened, why it happened, and how it happened, and then identify steps that can be taken to prevent a recurrence. This effort also might involve identifying who, if anyone, was or still is involved in the incident or attack.

56.5.6 Record. As discussed in Section 56.4.5 above, all CSIRT members should be trained to document everything during an incident. No relevant event or detail is too small to record when responding to computer emergencies. Always try to answer

"Who? What? Where? How? When? Why?" This is especially true when dealing with criminal activity, when there is the possibility or even an expectation that the intruder will be prosecuted. Keeping accurate records of what happened and the team's actions can prove pivotal in the organization's ability to identify positively the cause or source of an incident and to prevent similar incidents in the future.

In the case of criminal activity, with an expectation of prosecution, a legal representative should be kept informed so that appropriate forensic measures may be ordered at appropriate times and to avoid contaminating evidence through improper collection, storage, and analysis.

56.5.7 Return to Operations. For most business managers and executives, restoring operations is of paramount importance. Frequently, they will pressure the EDP people and the CSIRT to put off all other activities and to direct all resources to that end. Except in extreme cases, that pressure should be resisted, and the orderly carrying out of all preceding steps must be assured. As soon as possible, the CSIRT should assist operations personnel with bringing systems back online and returning them to full operating capacity. In some cases, hard drives, logs, and even entire systems may need to remain off-line until detailed forensics examinations can be completed. In these situations, backup systems should be used to bring systems and operating capabilities back online.

See Chapter 57 in this *Handbook* for a discussion of backups and Chapters 58 and 59 for extensive information about business continuity and disaster recovery.

56.5.8 Document and Review. While all CSIRT members should keep careful notes at all times, it is important to remember that formal procedures to document incidents and resulting actions are vital to the overall success of the incident response effort. This documentation can form the basis for new approaches, procedures, policies, awareness programs, and similar changes. Documenting successes and failures can provide the organization with a realistic view of its security posture and of its capability to respond to emergencies, and in some cases can justify the expenditure of additional funds on training or technology. This effort also ensures that the data captured can be used by the CSIRT to learn and to sharpen skills.

56.5.9 Involving Law Enforcement. The decision to involve law enforcement, or even when to involve law enforcement, in an incident response is one that must be given careful consideration. Although most organizations recognize the benefit of a close relationship with law enforcement, involving such agencies when responding to a computer emergency can have consequences beyond those that are immediately evident.

Clearly, local and national law enforcement agencies have a great deal to offer when establishing CSIRTs and developing incident handling capabilities. It is not uncommon these days to find that many law enforcement agencies have specialized units dedicated to computer crimes and issues. They can be a valuable resource, likely to have a wealth of threat data on hand. For this reason, it is important to partner with appropriate agencies to take advantage of their experience and to establish relationships. Knowing whom to contact in an emergency not only will save time and frustration, but may mean the difference between merely repelling an attack or catching and successfully prosecuting the perpetrator, which could help prevent future attacks.

Local laws and statutes may dictate specific notification requirements that an organization is obliged to follow in the event an actual or suspected incident occurs. Careful

review of local laws, statutes, and ordinances should be undertaken to ensure that the organization complies with notification requirements and other legal requirements.

When there is a choice to be made, the organization must weigh carefully the decision to involve law enforcement, and especially the question of when to do so. In most cases, formally involving law enforcement means that the organization may have to turn control of the incident and subsequent investigation over to the agency whose jurisdiction it is to investigate the crime.

While most professional law enforcement agencies will work with an organization to minimize any adverse impact on normal operations, this may not always be feasible. Because the goals of law enforcement often are different from those of others, especially of commercial enterprises, law enforcement agencies may not consider the impact of their response on the organization under attack.

Since law enforcement's mission is to investigate criminal activity, its focus will naturally be on identifying, tracking, and locating the intruder. This can, in some cases, result in seizure and removal for forensic purposes of systems and data, even systems that may be critical to the continued operation of the organization. In the case of a business, this might well mean the loss of necessary servers or workstations while an investigation is under way, with a possibly devastating effect.

Indeed, decisions about a preferred response may be taken out of the hands of managers and executives when law enforcement enters into an incident response situation. A commercial business might focus on identifying the vulnerability that made the attack possible, protecting against that vulnerability, and restoring systems to full operating capability. If, during the course of these efforts, the perpetrator can be identified, law enforcement will be informed, but such identification is rarely the overriding objective of the business. For law enforcement, however, identification and prosecution of the perpetrator is the primary objective. Establishing contact with appropriate law enforcement agencies before the organization is forced to respond to an incident will help the CSIRT plan when to notify law enforcement and how most effectively to align both sets of objectives when dealing with an incident.

For more details of working with law enforcement agencies and personnel, see Chapter 61 in this *Handbook*. For additional discussion of data collection for forensic applications, see Section 56.6.9 in this chapter.

56.5.10 Need to Know. Protecting information about an incident in progress is essential, not only to a successful response but because it can have serious legal, privacy, and other ramifications as well. Those charged with handling an incident must use out-of-band communications, such as cellular telephones, pagers, and encrypted email systems not connected to the system under attack, to ensure that knowledge of the incident is restricted to those who need to know about it. Attackers could intercept team communications if passed through in-band or normal channels and use that information to cover their tracks or even to prolong an incident.

Responding to incidents always involves gathering information about systems, users, activities, and events. In most cases, sensitive system and even personal information may be collected. During the course of their response and investigation, members of the team frequently make assumptions about the identities of those responsible for the incident, These assumptions are based on data that are continually being collected, refined, modified, and frequently changed during the course of an emergency response. Should an unproven or interim assumption that a particular individual was involved in the incident be made public, that individual's reputation might become needlessly tarnished, and the organization might well find itself facing legal proceedings as a result.

It is therefore essential that the CSIRT disseminate information about the incident according to a strict need-to-know policy. Limiting knowledge about an incident will help ensure that sensitive information remains in the hands of those who need it to perform their duties.

56.6 MANAGING THE CSIRT. All the work that goes into creating a CSIRT can be wasted if managers fail to lead. Sloppy management can result in degraded performance, alienation of the client base, staff frustration, sabotage, and employee turnover. Management plays a key role in the formation, operation, and support of a CSIRT. Ideally, teams should be composed not only of technical personnel, but also of managers with sufficient authority to assist the team in taking actions that contain an incident and that protect data and systems from further compromise. Outside of incident handling, management support for planning, establishing and enforcing policies, and preauthorizing responses is essential. Most important is management support of the CSIRT. Without solid backing from the highest levels of management, the CSIRT will be frustrated in its attempts to carry out its mission.

56.6.1 Professionalism. The DISA course wisely emphasizes the importance of professional behavior by all members of the CSIRT. The authors write:

> The survival of your CIRT may well depend upon using a Code of Conduct, which will earn the trust and respect of the commands you support. The conduct of any single team member reflects upon the entire CIRT organization. If the commands don't trust your CIRT, they won't report to you. It is important, therefore, not only to have a Code of Conduct, but also to shake it out and dust it off every once in a while. Remind team members what it is and why it is important… and use it.[15]

Here are some of the practical recommendations from that course (although we have put them in our own words for the most part):

- Write down the rules—a code of conduct—that represent your ideals of courteous, professional service to your clients.
- Train the team to understand and apply the code.
- Review the code periodically with the team.
- Speak clearly and avoid technobabble.
- Tell people exactly what you intend to do.
- Never hesitate to say "I don't know—but I'll find out."
- Do not criticize other people in your interactions with clients.
- Respect the confidentiality of your clients.
- Be respectful of your callers; do not belittle them or make them feel bad.

Notice how consistently we refer to clients; this usage emphasizes that both technical support teams and CSIRTs all perceive users as people to whom we owe service. There is no benefit to allowing an adversarial relationship between the technical support team or a CSIRT and the client base. Managers must not allow a gulf to develop between the CSIRT and the client community; leaders should clamp down on disparaging terms and derogatory comments about users. Team members must understand why such language is harmful.

Managers should identify CSIRT members with a chip on their shoulder; they must not adopt defensive, arrogant, or aggressive attitudes toward the users. If a computer security incident can be traced to procedural errors (i.e., the procedures themselves rather than user error are causing problems), the person reporting the problem should be thanked for the information, not criticized for having experienced or identified the problem.

No one in a CSIRT has ever regretted being professional.

56.6.2 Setting the Rules for Triage. As we mentioned in Section 56.4.2, *triage* in French means *sorting*. The same concept has been applied to helpdesks. For example, the "Helpdesk triage policy" from Courtesy Computers illustrates how a helpdesk team can categorize problems to ensure that important issues receive faster service than less important problems.[16] Importance is defined in terms of the number of users affected, the effects on mission-critical functions, and the costs of downtime or of less-than-optimal functions. The five priority levels suggested in the document are typical of the kind of triage categories established in many helpdesk departments (adapted from a table in the Courtesy Computers document):

Priority 1

- Issues of the highest importance; mission-critical systems with a direct impact on the organization (Examples: widespread network outage, payroll system, sales system, telecom system, etc.)
- Contact: Immediate–5 minutes
- Resolution: 30 minutes

Priority 2

- Single-user or group outage that is preventing the affected user(s) from working (Examples: failed hard drive, broken monitor, continuous OS lockups, etc.)
- Contact: 15 minutes
- Resolution: 1 hour

Priority 3

- Single-user or group outage that can be permanently or temporarily solved with a workaround (Examples: malfunctioning printer, synchronization problem, PC sound problem, etc.)
- Contact: 30 minutes
- Resolution: Same day

Priority 4

- Scheduled work (Examples: new workstation installation, new equipment/software order, new hardware/software installation)
- Contact: 1 hour
- Resolution: 1–4 days

Priority 5

- Nonessential scheduled work (Examples: office moves, telephone moves, equipment loaners, scheduled events)

- Contact: Same day
- Resolution: 5 days

The particular structure and specific timelines are merely examples, not blanket recommendations. Every organization must determine its own version of such a table.

In his helpful overview, "CIRT—Framework and Models," Ajoy Kumar summarizes the functions of triage in this way:

> Triage: The actions taken to categorize, prioritize, and assign incidents and events. It includes the following sub-processes:

- Categorize events.
- Correlate various events. Personnel involved in such teams typically also belong to Forensic teams.
- Prioritize events.
- Assign events for handling and response.
- Communicate information to "Respond" process for further handling.
- Re-assign (and close) events not belonging to CIRT.[17]

The DISA training materials suggest three broader categories of interactions with helpdesks and CSIRTs: "incidents, vulnerabilities, and information requests."[18]

1. *Incidents* involve breaches of security
2. *Vulnerabilities* include reports of security weaknesses (and may be reported as part of an incident)
3. *Information requests* often are managed using lists of frequently asked questions (FAQs).

The DISA instructors go on to define factors that can help CSIRTs prioritize incidents:

- The sensitivity and/or criticality of the data affected
- The amount of data affected
- Which host machines are involved
- Where and under what conditions the incident occurred
- Effects of the incident on mission accomplishment
- Whether the incident is likely to result in media coverage
- Number of users affected
- Possible relationships to other incidents currently being investigated
- The nature of the attack
- Economic impact and time lost
- Number of times the problem has recurred
- Who reports the incident[19]

On this last point, the DISA writers point out that the organizational rank of someone calling in an incident *may* bear on its priority—but that it may be wise to cross-check the report with a security expert who can speak to whether the report is sound. Sometimes a high-level manager's sense of urgency may be rooted more in his sense of self-importance than in operational requirements.

In summary, it is important to establish a sound basis for staff members of the CSIRT to carry out triage effectively. Once the rules for evaluating incidents have been clarified, staff members should practice analyzing a number of cases to train themselves in applying the rules consistently. Role-playing exercises, based on historical records or on made-up examples, can provide an excellent and enjoyable mechanism for staff members to establish a common standard for this difficult and sensitive task.

For additional ideas on using ideas and lessons from social psychology in managing security personnel, see Chapter 50 in this *Handbook*.

56.6.3 Triage, Process, and Social Engineering. Sometimes staff (or even managers) question the value of strict adherence to policy. Policy is sometimes seen as the expression of unnecessary rigidity—an inability to respond quickly to changing or unexpected circumstances. However, in CSIRT management, knowing and adhering to well-thought-out policies and following a reliable process are particularly valuable, not only for information gathering, data recording, and analysis, but also to maintain strict security.

One of the well-known tricks used by criminal hackers and spies is to simulate urgency that supports demands for violations of normal security restrictions. For example, criminals will call a relatively low-status employee, such as a secretary, and pressure him into violating standard protocols to obtain the password of his boss by claiming extreme circumstances of great urgency. The criminal may escalate the pressure to outright bullying by threatening the employee with punishment.

A criminal determined to penetrate security barriers can manufacture an incident that leads to involvement of the CSIRT. Allowing such a person to apply pressure for violations of protocol is an invitation to compromise. Worse, such deviations from well-tried and well-justified procedures can add to the embarrassment caused by the compromise; it is bad enough to have someone breaking through our security without having to admit that we helped.

For additional information on social engineering, see Chapters 19 and 20 in this *Handbook*.

56.6.4 Avoiding Burnout. Much of the discussion that follows applies equally to CSIRTs and to helpdesks; in a sense, one can view the CSIRT as a specialized helpdesk. Many CSIRTs are specialized subsets of the helpdesk team.

Any organization, even one with a relatively small CSIRT or a small helpdesk, can suffer spikes in demand. Ordinary business cycles can influence network usage; for example, universities often see perfectly normal but large increases in call volumes at registration times as new students forget their passwords, try to connect unverified laptops to the university network, or get blocked for violating appropriate-use policies. At any site, a denial-of-service attack, a plague of computer virus infections, or an infestation of computer worms can cause a flood of calls.

Another trend is the ironic observation that the better a CSIRT (or helpdesk team, although the focus will continue on CSIRTs) becomes at handling problems, the more readily members of its community will turn to it to report problems or ask for help.

Thus, the better the CSIRT does its job, the heavier its workload can become, at least for a while. According to the DISA course:

> As a new CSIRT grows and the workload increases, and especially on those teams that provide 24-hour emergency response, burnout becomes quite common. By studying the issue, one national CSIRT determined that a full-time team member could comfortably handle one new incident per day, with 20 incidents still open and actively being investigated.[20]

Staff members who face increasing workloads may become stressed. Working long periods of overtime, missing time with family and friends, perhaps even missing regular exercise and food—these factors may lead to increased errors and turnover if people are forced to accept increasingly demanding conditions for long periods.

One of the most valuable organizational approaches to preventing burnout is to rotate staff from the IT group through the CSIRT function on a predictable schedule. For example, one can assign people to the CSIRT for three- or six-month rotations.

Such rotations require especially good training programs and particularly good documentation, to maintain efficiency as new people come on duty; in addition, the assignments must be staggered so that the CSIRT does not have to cope with large numbers of newcomers all at once. Ideally, there would not be more than one switch of personnel a week.

How should existing assignments be transferred within the CSIRT? *Difficult* existing cases should be transferred to staff members who have been on duty for a few weeks, not to the incoming staff member (even one with experience on the CSIRT). The incoming CSIRT member should be given a chance to get into (or get back into) the rhythm of the job before being hit with the most intractable problem or the most ornery client.

Every incident must have a case coordinator—the person who monitors the problem, aggregates information from varied resources, and serves as the voice of the CSIRT for that incident. When transferring responsibility for a case from one case coordinator to another, managers should ensure that the previous coordinator prepares clients for the transition and introduces the new coordinator to the key client contacts to ensure a smooth transition of control. Clients often come to depend on the person they have been working with to resolve an incident; an unexpected change can be unsettling and disturbing.

56.6.5 Many Types of Productive Work. The DISA course writers suggest:

> Allow team members to allocate time away from high stress incident response assignments and pursue broader interests in areas such as tool development, public education and presentations, research, and other professional opportunities.[21]

CSIRT members, by the nature of their work, will have a great deal to contribute to the awareness, training, and education of their colleagues.

In the technical support group for Hewlett-Packard Canada in the 1980s, managers exercised great care in preventing consistent overwork. In emergencies, employees all pitched in—including managers—to resolve the problem for clients; however, the policy on allocation of time was strictly enforced under normal circumstances. Everyone kept careful records of time worked—a habit everyone can usefully follow—so that managers could analyze where the burden was distributed, and to provide statistical information for load balancing and personnel planning. Employees who violated the

policy that no more than 70 percent of their time in the system engineering group should be spent on billable hours were warned to ease off; the rationale was that it was *necessary* to maintain constant training and time for administration and just for thinking, to ensure long-term productivity of their specialists.

56.6.6 Setting an Example. The behavior of managers can greatly influence morale, motivation, and dedication among team members. For example, supervisors and upper managers can greatly motivate staff by pitching in to support them during emergencies or extraordinary demands, even if only by their presence. Making the CSIRT a stimulating and enjoyable duty that people *want* to be on is one of the best approaches to avoiding burnout and ensuring reliable response to computer-related problems.

56.6.7 Notes on Shiftwork. As discussed in Section 56.6.4, rotating assignments among CSIRT members can be an excellent idea. However, frequent changes in work schedules that involve changes in sleep cycles are *not* a good idea; for example, weekly changes in shift from day to night schedules can seriously disrupt the natural circadian wake/sleep cycle and have been shown to increase the rate of errors and accidents.[22] One authoritative resource states that there are "adverse health and safety effects to working shifts[:]"

> A shiftworker, particularly one who works nights, must function on a schedule that is not natural. Constantly changing schedules can:
>
> - upset one's circadian rhythm (24-hour body cycle),
> - cause sleep deprivation and disorders of the gastrointestinal and cardiovascular systems,
> - make existing disorders worse, and
> - disrupt family and social life.[23]

Scientific studies throughout the world have long shown that shiftwork, by its very nature, is a major factor in the health and safety of workers; LaDou writes in his abstract:

> Daily physiologic variations termed *circadian rhythms* are interactive and require a high degree of phase relationship to produce subjective feelings of wellbeing. Disturbance of these activities, circadian desynchronization, whether from passage over time zones or from shift rotation, results in health effects such as disturbance of the quantity and quality of sleep, disturbance of gastrointestinal and other organ system activities, and aggravation of diseases such as diabetes mellitus, epilepsy, and thyrotoxicosis.[24]

The U.S. National Institute for Occupational Safety and Health has published a monograph about shiftwork that contains this advice for improving shiftwork schedules:

- Avoid permanent (fixed or nonrotating) night shift
- Keep consecutive night shifts to a minimum
- Avoid frequent shift changes—provide enough stability to let employees adapt to their schedule
- Plan some free weekends

- Avoid several long days of work followed by four- to seven-day "mini-vacations"—such schedules are stressful because of the radical shifts of diurnal cycles in the two phases
- Keep long work shifts and overtime to a minimum
- Consider different lengths for shifts
- Examine start-end times to fit in better with life in the external world; for example, making it possible for parents on night shift to see their children before they leave for school
- Keep the schedule regular and predictable
- Examine rest breaks[25]

56.6.8 Role of Public Affairs. The nature of interconnected systems today—as well as some employees' practice of blurting out confidential information on their social-networking pages—all but guarantees that any incident will become obvious to partners, customers, clients, and others, including even the press. In many cases, the organization will be compelled to advise its constituents continuously of the status of any outage or degradation of services resulting from an incident, and the causes behind it. Therefore, it is crucial that information released for general consumption be properly screened and cleared prior to release.

It is equally important that such information be released through a single source, such as the public affairs office. Restricting release of incident-related information through the public affairs office, or other designated point, will help ensure that frequent, straightforward communications with stakeholders can take place, while at the same time controlling rumors and misinformation. This simple step can do much to lessen anxiety about an incident and to reassure members, partners, and customers that the situation is well in hand and will be resolved.

56.6.9 Importance of Forensic Awareness. As implied in Section 56.5.9, depending on the specific incident, the organization may desire not only to control the incident but also to trace and prosecute the perpetrators in the case of known or suspected criminal activity. It is therefore highly advisable that members of the CSIRT receive thorough training in procedures for collecting and preserving evidence. Mishandling of evidence can result in an inability to take successful legal action against an attacker or to recover damages following an incident. Computer forensics and evidence handling should be high on the CSIRT's list of training topics. Chapters 2 and 61 of this *Handbook* contain additional material on computer forensics and working with law enforcement.

56.7 POSTINCIDENT ACTIVITIES. One of the most important principles of management in general, and operations management in particular, is that fixing a problem has two aspects: the short term and the long term. One must be able to solve problems quickly enough to be effective; that is, the speed of solution must be appropriate to the consequential costs of delay. However, we should not figuratively wipe our hands in satisfaction and walk away from a problem resolution without thinking about why it happened, how we fixed it, and whether we can do better to avoid repeats and to improve our response.[26]

The CSIRT's efforts do not end once the incident is resolved. Instead, the team should take a reasonable period to rest and recover. Then, while the details are still

fresh in the team members' minds, they should examine the incident from start to finish, both formally and informally, asking questions such as "What happened? What went right? What went wrong?" This way, the team will learn from each incident and become more efficient and confident when handling new incidents in the future.

At the conclusion of each incident, the team should be assembled, and a formal debriefing and review of the incident should be carried out. This debriefing should include a complete review of the team and its handling of the incident, including its adherence to policy and its technical performance. Each team member should be individually debriefed following the incident. Their recollections, thoughts, ideas, and reactions as to how the incident was handled and how the team performed, should be documented and preserved. A management team might debrief members, and team members might debrief each other, or they might even debrief themselves using a checklist or form. Regardless of the method, the CSIRT members themselves are the best source of data about the weaknesses and strengths of the team, and that data must be captured if the team is to improve and grow in skills and confidence.

Once individual impressions are captured, it is often effective to assemble the team as a group for an incident postmortem session. Starting from the beginning of the incident, the team should examine whether it had adequate, workable policies on which to base its actions and decisions. The group should jointly evaluate each aspect of the team, its composition, skills, authority, and step-by-step handling of the incident. A list of lessons learned and action items for improvements should result from this review.

Data collected during this review process should form the basis for improving the team. This information provides input to what should be a continuous cycle involving planning, preparation, training, responding, and evaluating. Shortfalls in training, skills, equipment, access, policies, and authority will become evident through this process. These shortfalls can be corrected to improve the team's ability to respond effectively to incidents in the future.

The next sections provide additional insights into how to learn from the CSIRT's experiences.

56.7.1 Postmortem. As a matter of standard operating procedure, every technical support person and the CSIRT must schedule time to analyze the underlying factors that led to the problem they have just resolved. This analysis will likely involve operational staff outside the CSIRT; these are the people with line expertise who will be able to contribute their intimate knowledge of technical details that contributed to this security breach. These discussions can often lead to practical recommendations for improvement of the security architecture, such as its topology or firewall placement, operational procedures such as monitoring standards or vulnerability patching, and technical details such as configurations or parameter settings.

Similarly, it is a commonplace in discussions of disaster recovery and business continuity planning that every practice run, or real-life incident, should be analyzed to see where we have made errors or achieved less than our goals in performance. Managers must ensure that these analyses are not perceived as (or worse, really are) finger-pointing exercises for apportioning blame. In a column for *Network World*, M. E. Kabay has explained the concepts of *egoless work*; the postmortem analysis of an incident must be ego-free.[27] Managers can set the tone by responding positively to what might otherwise be perceived as criticism; "That's a good point" and "Very good observation" are examples of positive, encouraging responses to observations such as "We were too slow in getting back to the initial caller given that she clearly stated that

the entire department was off-line." The meeting should focus on ways to improve the response, given the insights resulting from detailed analysis of successes and failures during an incident.

The other aspect that sometimes gets lost in such postmortems is exploring the reasons for the problems. If we do not pay attention to underlying causes, we may fix specific problems, and we may improve particular procedures, but we will likely encounter different consequences of the same fundamental errors that caused those particular problems. We must pursue the analysis deeply in order to identify structural flaws in our processes, so that we can correct those problems and thus reduce the likelihood of entire classes of problems.

The U.S. National Institute of Standards and Technology *Computer Security Incident Handling Guide* specifically recommends a postincident analysis in Section 3.4. We quote the authors' list of suggested questions verbatim:

- Exactly what happened, and at what times?
- How well did staff and management perform in dealing with the incident? Were the documented procedures followed? Were they adequate?
- What information was needed sooner?
- Were any steps or actions taken that might have inhibited the recovery?
- What would the staff and management do differently the next time a similar incident occurs?
- What corrective actions can prevent similar incidents in the future?
- What additional tools or resources are needed to detect, analyze, and mitigate future incidents?[28]

The authors also recommend these actions (paraphrasing and summarizing):

- Invite people to the postmortem with an eye to increasing cooperation throughout the organization.
- Plan the agenda by polling participants before the meeting.
- Use experienced moderators.
- Be sure the meeting rules are clear to everyone to avoid confusion and conflict.
- Keep a written record of the discussions, conclusions, and action items.

On this last point, we add that all action items should indicate clearly who intends to deliver precisely what operational result, to whom, in which form, and by when.

56.7.2 Continuous Process Improvement: Sharing Knowledge within the Organization. On page 3-23 of the *Computer Security Incident Handling Guide*, the authors make a series of recommendations on how to capitalize on the knowledge gained through systematic analysis of incidents.[29] We are commenting briefly on each of their suggestions (which are shown in quotation marks).

- "Reports from these meetings are good material for training new team members by showing them how more experienced team members respond to incidents." The incident reports that were used for discussion in the analytic meetings should

be made available, perhaps as appendices, in a single report document so that all of the information about a specific incident or series of incidents can be accessed at one time. In what follows, such a dossier is referred to as the *follow-up report*.

- "Another important postincident activity is creating a follow-up report for each incident, which can be quite valuable for future use." The general principle is that without documentation, we lose the opportunity for increasing institutional knowledge. If we do not record what we have learned, transmission depends on luck: the haphazard contacts of people who need to know something with those who can help. Without documentation and efficient indexing, information transferred becomes an inefficient, random process of querying and guesswork. Informal knowledge sometimes remains limited to a few people or even a single individual; without these key resources, the information is unavailable. If the holders of undocumented information leave the organization, their knowledge is usually lost to the group.
- "First, the report provides a reference that can be used to assist in handling similar incidents." Why waste time reinventing solutions that have already been found? Why make the same errors and cause the same problems that have already been located and that could be avoided?
- "Creating a formal chronology of events (including time-stamped information such as log data from systems) is important for legal reasons, as is creating a monetary estimate of the amount of damage the incident caused in terms of any loss of software and files, hardware damage, and staffing costs (including restoring services)." One of the most important kinds of information for managing security is the cost estimate. Rational allocation of resources depends on knowing how often problems occur, and how much they cost, so that we can spend appropriate amounts of money for equipment, and for the time of our employees and consultants to prevent such problems.
- "This estimate may become the basis for subsequent prosecution activity by entities such as the U.S. Attorney General's office." Estimates of monetary consequences are also essential for civil torts in the calculation of restitution.
- "Follow-up reports should be kept for a period of time as specified in record retention policies." As the *Guide*'s authors discuss in their section 3.4.2, historical records become increasingly useful as they provide a statistical base for analyzing and predicting phenomena. The costs of saving such data (which have relatively small volumes) have dropped to virtually nothing given the huge digital storage capacities of today's archival media, and their extremely low cost.

56.7.3 Sharing Knowledge with the Security Community. One of the most valuable contributions we can make to each other is information sharing. The Computer Emergency Response Team Coordination Center (CERT/CC) offers an overview of why and how to report security incidents in its *Handbook for Computer Security Incident Response Teams*.[30] The CSIRT experts summarize the types of activity on which they would appreciate receiving reports; reasons for reporting security incidents; the variety of people and agencies who can benefit from such reports; extensive guidelines on what to include in the reports; and how to reach the CERT/CC securely.

The section "Why should I report an incident?" has these headers (and a paragraph or so of explanation of each point):

- You may receive technical assistance.
- We may be able to associate activity with other incidents.
- Your report will allow us to provide better incident statistics.
- Contacting others raises security awareness.
- Your report helps us to provide you with better documents.
- Your organization's policies may require you to report the activity.
- Reporting incidents is part of being a responsible site on the Internet.

Another way of contributing to the field is to speak at conferences. For example, the Forum of Incident Response and Security Teams (FIRST) organizes conferences, technical colloquia, and workshops.[31] The conferences are open to all, not just members of FIRST, and organizers want participants to:

- Learn the latest security strategies in incident management
- Increase their knowledge and technical insight about security problems and their solutions
- Keep up-to-date with the latest incident response and prevention techniques
- Gain insight on analyzing network vulnerabilities
- Hear how the industry experts manage their security issues
- Interact and network with colleagues from around the world to exchange ideas and advice on incident management best practices

Readers should think about contributing papers to such conferences. Anyone who has spoken at technical conferences will confirm that there is no better way to solidify one's expertise than marshaling information into a clear presentation and speaking before one's peers. Feedback from interested participants can improve not only the current presentation but also the process being described. Intelligent, enthusiastic interchange among practitioners of goodwill with varied experiences, and from different environments, is not only productive of new ideas, it is immense fun.

The FIRST event includes "Lightning Talks" which are described as "short presentations or speeches by any attendee on any topic, which can be scheduled into conference proceedings with the approval of the organisers." Participants with hot news can thus present their findings or their ideas without necessarily having to prepare a long lecture or submitting their work many months in advance.

Other conferences, such as those organized by the Computer Security Institute (CSI),[32] MIS Training Institute (MISTI),[33] and RSA Security,[34] among many others, usually offer opportunities for discussions of CSIRT management. Readers, if they can, should take advantage of these opportunities by registering for the calls for participation (CFPs) and responding to one or two a year.

56.8 CONCLUDING REMARKS. CSIRTs are an effective organizational tool for responding to computer emergencies. However, to be effective, these teams must be carefully planned, built, trained, and supported. Proper planning and the establishment of a clear set of organizational objectives for the CSIRT are key to ensuring success.

Teams that are well planned, well trained, confident, and that possess the authority and training to execute their stated mission, ultimately can provide a real return on investment for an organization. This return often can be measured in terms of limiting the impact and cost, both tangible and intangible, of a computer emergency.

56.9 FURTHER READING

Brownlee, N., and E. Guttman. "Expectations for Computer Security Incident Response," RFC 2350. IETF (June 1998), http://datatracker.ietf.org/doc/rfc2350/

Beisse, F. *A Guide to Computer User Support for Help Desk and Support Specialists*, 5th edition. Cengage Learning, 2012.

Czegel, B. *Running an Effective Help Desk*, 2nd ed. Wiley, 1998.

Gerbyshak, P., and J. M. Brooks. *Help Desk Manager's Crash Course*. BookSurge Publishing, 2009.

Johnson, L. R. *Computer Incident Response and Forensics Team Management: Conducting a Successful Incident Response*. Syngress, 2013.

Knapp, D. *A Guide to Service Desk Concepts*, 4th ed. Cengage Learning, 2013.

Lucas, J., and B. Moeller. *The Effective Incident Response Team*. Addison-Wesley Professional, 2003.

Mandia, K., C. Prosise, and M. Pepe. *Incident Response and Computer Forensics*, 2nd ed. McGraw-Hill/Osborne, 2003.

McCarthy, N. K., M. Todd, and J. Klaben. *The Computer Incident Response Planning Handbook: Executable Plans for Protecting Information at Risk*. McGraw-Hill Osborne Media, 2012.

Menken, I. *Transform and Grow Your Help Desk into a Service Desk within Service Operation: Service Desk, Help Desk Best Practice within Service Operation*. Emereo, 2010.

Rajnovic, D. *Computer Incident Response and Product Security*. Cisco Press, 2010.

Robertson, K. E. *Spectacular Support Centers*. Customer Service Press, 2007.

Online resources and some software links:

- BMC TrackIt!: www.numarasoftware.com/track-it
- CERT/CC. "Resources for Computer Security Incident Response Teams (CSIRTs)," 2012, www.cert.org/csirts/resources.html
- Help Desk Institute: www.thinkhdi.com
- HelpMaster Pro Suite: www.helpmasterpro.com
- Hesk Help Desk Software: www.hesk.com
- Open Source Ticket Request System (OTRS): http://otrs.org
- Request Tracker (RT): www.bestpractical.com/rt
- Zendesk: www.zendesk.com/help-desk-software

56.10 NOTES

1. Parts of this chapter are based on a long series of articles published over several years by M. E. Kabay in his *Network World Security Strategies* newsletter. To avoid cluttering the text with unnecessary endnotes and quotation marks, there are no further specific references to particular sources in the *Network World* series on CSIRT management. For an archive of Kabay's *Network World* newsletters in PDF, including a global index, see www.mekabay.com/nwss

2. T. Maleshefski, "Infographic: The History of Customer Support," www.zendesk.com/blog/the-history-of-customer-support
3. Computer Security Resource Center, "History of the Federal Computer Incident Response Capability (FedCIRC) Pilot," 1998, http://csrc.nist.gov/topics/incidentNIST/history.htm
4. www.cert.org
5. M. J. West-Brown, D. Stikvoort, K.-P. Kossakowski, G. Killcrece, R. Ruefle, and M. Zajicek, *Handbook for Computer Security Incident Response Teams (CSIRTs),* 2nd ed. (Computer Emergency Response Team Coordination Center [CERT/CC™], Carnegie Mellon University Software Engineering Institute, 2003); www.cert.org/archive/pdf/csirt-handbook.pdf
6. N. Brownlee and E. Guttman, "Expectations for Computer Security Incident Response," Internet Engineering Task Force Network Working Group Request for Comment (RFC) 2350 (June 1998), http://tools.ietf.org/html/rfc2350
7. DISA, "Introduction to Computer Incident Response Team (CSIRT) Management, v1.0" (CD-ROM), Defense Information Systems Agency, 2001. Available free for download as ZIP file with permission of the DISA Information Assurance Support Environment at www.mekabay.com/infosecmgmt/disa_cirtm_cdrom.zip
8. National Council of ISACs (Information Sharing and Analysis Centers) www.isaccouncil.org
9. Microsoft® Encarta® Reference Library 2008.
10. DISA, "Introduction to Computer Incident Response Team (CSIRT) Management, v1.0."
11. B. Czegel, *Running an Effective Help Desk*, 2nd ed. (Wiley, 1998).
12. J. D. Howard, "An Analysis of Security Incidents on the Internet, 1989–1995," Ph.D. diss., Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, April 1997, www.cert.org/archive/pdf/JHThesis.pdf
13. For an extensive list of articles on this topic, use the Network World search function (www.networkworld.com/search) and enter the term "help desk."
14. M. E. Kabay, "The Art of Technical Support," www.mekabay.com/courses/academic/jac/TSP/index.htm
15. DISA, "Introduction to Computer Incident Response Team (CSIRT) Management, v1.0."
16. "Help Desk Triage Policy," Courtesy Care Services, www.courtesycare.net/Best%20Practices/help%20desk%20triage.doc
17. A. Kumar, "CIRT—Framework and Models." SecurityDocs.com, January 31, 2005, originally available at www.securitydocs.com/library/2964 and now available at www.docstoc.com/docs/84310604/CIRT—Framework-and-Models
18. DISA, "Introduction to Computer Incident Response Team (CSIRT) Management, v1.0."
19. DISA, "Introduction to Computer Incident Response Team (CSIRT) Management, v1.0."
20. DISA, "Introduction to Computer Incident Response Team (CSIRT) Management, v1.0."
21. DISA, "Introduction to Computer Incident Response Team (CSIRT) Management, v1.0."

22. A. Kerin and A. Aguirre, "Improving Health, Safety, and Profits in Extended Hours Operations (Shiftwork)," *Industrial Health* 43 (2005):201–208, www.jniosh.go.jp/en/indu_hel/pdf/43-1-28.pdf
23. Canadian Centre for Occupational Health and Safety, "Rotational Shiftwork," 2010, www.ccohs.ca/oshanswers/work_schedules/shiftwrk.html
24. J. LaDou, "Health Effects of Shift Work," *Western Journal of Medicine* 137, No. 6 (December 1982), www.ncbi.nlm.nih.gov/pmc/articles/PMC1274227
25. R. R. Rosa, "Plain Language about Shiftwork," U.S. Department of Health and Human Services, Public Health Service, Centers for Disease Control and Prevention, National Institute for Occupational Safety and Health, 2001. Available as a NOOK document from Barnes & Noble, www.barnesandnoble.com/w/plain-language-about-shiftwork-roger-rosa/1004120554?ean=9780160615603
26. M. E. Kabay, "On Not Knowing: Embracing the Opportunity to Learn," 2004, www.mekabay.com/opinion/not_knowing.pdf
27. M. E. Kabay, "Egoless Work: Take Your Ego out of the Equation," *Network World Security Strategies Newsletter*, February 2, 2006, www.mekabay.com/nwss/435_egoless_work.pdf
28. P. Cichonski, T. Millar, T. Grance, and K. Scarfone, *Computer Security Incident Handling Guide*, NIST Special Publication SP800-61 Rev. 2, August 2012. http://csrc.nist.gov/publications/nistpubs/800-61rev2/SP800-61rev2.pdf
29. Cichonski et al., *Computer Security Incident Handling Guide.*
30. M. J. West-Brown, D. Stikvoort, K.-P. Kossakowski, G. Killcrece, R. Ruefle, and M. Zajicek, *Handbook for Computer Security Incident Response Teams (CSIRTs)*, 2nd ed.
31. www.first.org
32. www.gocsi.com
33. www.misti.com
34. www.rsaconference.com

CHAPTER 57

DATA BACKUPS AND ARCHIVES

M. E. Kabay and Don Holden

57.1 INTRODUCTION 57·1
- 57.1.1 Definitions 57·2
- 57.1.2 Need 57·2

57.2 MAKING BACKUPS 57·3
- 57.2.1 Fault-Tolerant Systems 57·3
- 57.2.2 Hierarchical Storage Systems 57·3
- 57.2.3 Disk Mirroring 57·3
- 57.2.4 Logging and Recovery 57·6
- 57.2.5 Backup Software 57·7
- 57.2.6 Removable Media 57·7
- 57.2.7 Labeling 57·13
- 57.2.8 Indexing and Archives 57·14
- 57.2.9 Registry Backups 57·15
- 57.2.10 Bootable Versions 57·15
- 57.2.11 Cloud Backup 57·16

57.3 BACKUP STRATEGIES 57·18
- 57.3.1 Exclusive Access 57·19
- 57.3.2 Types of Backups 57·19
- 57.3.3 Computer Systems 57·21
- 57.3.4 Testing 57·24
- 57.3.5 Known-Good Boot Media 57·24

57.4 DATA LIFE CYCLE MANAGEMENT 57·24
- 57.4.1 Retention Policies 57·25
- 57.4.2 Rotation 57·25
- 57.4.3 Media Longevity and Technology Changes 57·26

57.5 SAFEGUARDING BACKUPS 57·27
- 57.5.1 Environmental Protection 57·27
- 57.5.2 Onsite Protection 57·27
- 57.5.3 Offsite Protection 57·28

57.6 DISPOSAL 57·29
- 57.6.1 Scavenging 57·30
- 57.6.2 Data and Media Destruction 57·30

57.7 COSTS 57·32

57.8 OPTIMIZING FREQUENCY OF BACKUPS 57·33

57.9 CONCLUDING REMARKS 57·34

57.10 FURTHER READING 57·35

57.11 NOTES 57·35

57.1 INTRODUCTION. Nothing is perfect. Equipment breaks, people make mistakes, and data files become corrupted or disappear. Reliably dated backups can provide evidence of innocence (or guilt) in forensic examinations and for proof of legitimate authorship in copyright disputes. Everyone, and every system, needs a well-thought-out backup and retrieval policy. In addition to making backups, data processing personnel also must consider requirements for archival storage and for retrieval of data copies. Backups also apply to personnel, equipment, and

electrical power; for other applications of redundancy, see Chapters 23 and 45 in this *Handbook*.

57.1.1 Definitions. **Backups** are copies of data files or records, made at a moment in time, and primarily used in the event of failure of the active files. Normally, backups are stored on different media from the original data. In particular, a copy of a file on the same disk as the original is an acceptable backup only for a short time; the *.bak, *.bk!, *.wbk, and *.sav files created by programs such as word processors are examples of limited-use backups. However, even a copy on a separate disk loses value as a backup once the original file is modified, unless *incremental* or *differential* backups also are made. These terms are described in Section 57.3.3. Typically, backups are taken on a schedule that balances the costs and inconvenience of the process with the probable cost of reconstituting data that were modified after each backup; rational allocation of resources is discussed in Section 57.8.

Deletion or corruption of an original working file converts the most recent backup into the *de facto* original. Those who do not understand this relationship mistakenly believe that once they have a backup, they can safely delete the original file. However, before original files are deleted, as when a disk volume is to be formatted, there must be at least *double backups* of all required data. Double backups will usually ensure continued operations should there be a storage or retrieval problem on any one backup medium.

This chapter uses these **abbreviations** to denote *data storage capacities*:

KB = kilobyte = 1,024 bytes (characters) (approximately 10^3)
MB = megabyte = 1,024 KB = 1,048,576 bytes ($\sim 10^6$)
GB = gigabyte = 1,024 MB = 1,073,741,824 bytes ($\sim 10^9$)
TB = terabyte = 1,024 GB = 1,099,511,627,776 bytes ($\sim 10^{12}$)
PB = petabyte = 1,024 TB = 1,125,899,906,842,624 bytes ($\sim 10^{15}$)
EB = exabyte = 1,024 PB = 1,152,921,504,606, 846,976 bytes ($\sim 10^{18}$)[1]

Archives—Although backups are used to store and retrieve operational data chronologically, data archives are used to store data that is seldom or no longer used but, when needed for reference, compliance (e.g., Federal Drug Administration [FDA], Health Insurance Portability and Accountability Act of 1996 [HIPAA], Securities and Exchange Commission [SEC]), or legal actions such as e-discovery (turning over electronic evidence in a legal proceeding), must be retrieved in a quick, logical mode. Some new data storage management systems combine both backup and this expanded archival capability.

Cloud backups—Backup copies can also be kept on servers connected to the originating systems through the Internet; also known as *online storage*.

57.1.2 Need. Backups and storage archives are used for many purposes:

- To replace lost or corrupted data with valid versions
- To satisfy audit and legal requirements for access to retained data
- In forensic examination of data to recognize and characterize a crime and to identify suspects

- For statistical purposes in research
- To satisfy requirements of due care and diligence in safeguarding corporate assets
- To demonstrate the authorship of documents or works of art through a chain of documented, dated copies
- To meet unforeseen requirements

57.2 MAKING BACKUPS. Because data change at different rates in different applications, backups may be useful when made at frequencies ranging from milliseconds to years.

57.2.1 Fault-Tolerant Systems. The ultimate backup strategy is to do everything twice at the same time. Computer systems such as HP NonStop (formerly Tandem) and Stratus use redundant components at every level of processing; for example, they use arrays of processors, dual input/output (I/O) buses, multiple banks of random access memory, and duplicate disk storage devices to permit immediate recovery should anything go awry. Redundant systems use sophisticated communications between processors to ensure identity of results. If any computational components fail, processing can continue uninterrupted while the defective components are replaced.

A similar approach duplicates disk I/O, as discussed below in Section 57.2.3 on disk mirroring.

57.2.2 Hierarchical Storage Systems. Large computer systems with terabytes or petabytes of data typically use a *hierarchical storage system* to place often-used data on fast, relatively expensive disks while migrating less-used data to less expensive, somewhat slower storage media such as slower magnetic disks, optical media, or magnetic tapes. However, users need have no knowledge of, or involvement in, such migration; all files are listed by the file system and can be accessed without special commands. Because the secondary storage media are stored in dense cylindrical arrays, usually called *silos,* they may have total capacities in the PB per silo, with fast-moving robotic arms that can locate and load the right unit within seconds. Users may experience a brief delay of a few seconds as data are copied from the secondary storage units back onto the primary hard disks, but otherwise there is no problem for the users. This system provides a degree of backup simply because data are not erased from the silo media when they are copied to magnetic disk on the system, nor are data removed from disk when they are appended to the secondary storage; this data remanence provides a degree of temporary backup because of the duplication of data. For more information on disk-based storage systems, see Chapter 36 in this *Handbook*.

57.2.3 Disk Mirroring. There are several methods for duplicating disk operations so that disk failures cause limited or no damage to critical data. *Continuous data protection* (CDP) provides infinitely granular *recovery point objectives* (RPO), and some implementations can provide near instant *recovery time objectives* (RTOs). This is because CDP protection typically is done on a write-transaction-by-write-transaction basis. Every modification of data is recorded, and the recovery can occur to any point, down to the demarcation of individual write operations.

The Continuous Data Protection Special Interest Group (CDP SIG) of the Storage Networking Industry Association (SNIA) defines CDP as:

> …a methodology that continuously captures or tracks data modifications and stores changes independent of the primary data, enabling recovery points from any point in the past. CDP systems may be block-, file-, or application-based and can provide fine granularities of restorable objects to infinitely variable recovery points. So, according to this definition, all CDP solutions incorporate these three fundamental attributes:
>
> 1. Data changes are continuously captured or tracked
> 2. All data changes are stored in a separate location from the primary storage
> 3. Recovery point objectives are arbitrary and need not be defined in advance of the actual recovery[2]

57.2.3.1 RAID. *Redundant arrays of independent* (originally *inexpensive*) *disks* (RAID) were described in the late 1980s and have become a practical approach to providing fault-tolerant mass storage. The falling price of disk storage has allowed inexpensive disks to be combined into highly reliable units, containing different levels of redundancy among the components, for applications with requirements for full-time availability. The disk architecture involves special measures for ensuring that every sector of the disk can be checked for validity at every input and output operation. If the primary copy of a file shows data corruption, the secondary file is used and the system automatically makes corrections to resynchronize the primary file. From the user's point of view, there is no interruption in I/O and no error.

RAID levels define how data are distributed and replicated across multiple drives. The most commonly used levels are described in Exhibit 57.1.[3]

57.2.3.2 Storage Area Network and Network Area Storage. A storage area network (SAN) is a network that allows more than one server to communicate with more than one device via a serial Small Computer System Interface (SCSI) protocol. SANs share devices via the SCSI protocol, running on top of a serial protocol such as Fibre Channel or Internet SCSI (iSCSI). They offer multiple servers access to raw devices, such as disk and tape drives.

Network-Attached Storage (NAS) effectively moves storage out from behind the server and puts it directly on the transport network. Unlike file servers that have SCSI and local area network (LAN) adapters, an NAS appliance offers multiple servers access to a file or file system via standard file-sharing protocols: Network File System (NFS), Common Internet File System (CIFS), Direct Access File System (DAFS), or others. One NAS disadvantage is that it shifts storage transactions from parallel SCSI connections to the production network. Therefore, the backbone LAN has to handle both normal end user traffic and storage disk requests, including backup operations. For more information on SAN and NAS, see Chapter 36.

57.2.3.3 Workstation and Personal Computer Mirroring. Anyone who has more than one computer may have to synchronize files to make them the same on two or more computers. In addition to the simple matter of convenience, synchronizing the computers provides excellent backup, which supplements daily incremental backups (i.e., backups of all the files that have changed since the previous incremental backup) and periodic full backups. The inactive synchronized computers serve as daily full backups of the currently active computer. For example, FreeFileSync,

RAID Features and Performance

Comparison of RAID levels from the RAID Advisory Board.

Common Name	Description	Disks (cost)	Data Reliability	Data Transfer	Maximum I/O Rate
0 Disk Striping	Data is distributed across disks in the array. No redundant info provided.	*N*	Lower than single disk	Very high	Very high for read and write
1 Mirroring	All data replicated on *N* separate disks. *N* is almost always 2.	2*N* 3*N* etc.	Higher than RAID 2, 3, 4, or 5. lower than 6	R: higher than single disk W: similar to single disk	R: up to 2x single disk W: similar to single disk
2	Data are protected by Hamming code. Redundant information distibued across *m* disks (*m* = number of data disks in array).	*N*+*m*	Much higher than single disk; comparable to RAID 3, 4, or 5	Highest	Similar to 2x single disk
3 Parallel Transfer Disks with Parity	Each data sector is subdivided and distributed across all data disks. Redundant information normally stored on dedicated parity disk.	*N*+1	Much higher than single disk; comparable to RAID 2, 4, or 5	Highest	Similar to 2x single disk
4	Data sectors distributed as with disk striping. Redundant information stored on dedicated parity disk.	*N*+1	Much higher than single disk; comparable to RAID 2, 3, or 5	R: similar to disk striping W: much lower than single disk	R: similar to disk striping W: much lower than single disk
5	Data sectors distributed as with disk striping. Redundant information interspersed with user data.	*N*+1	Much higher than single disk; comparable to RAID 2, 3, or 4	R: similar to disk striping W: lower than single disk	R: similar to disk striping W: usually lower than single disk
6	As RAID Level 5, but with additional independently computed redundant information	*N*+2	Highest	R: similar to disk striping W: lower than RAID 5	R: similar to disk striping W: much lower than RAID 5

Exhibit 57.1 RAID Levels *Source*: Used with kind permission of Alan Freedman, Computer Desktop Encyclopedia (Computer Language Company, 2013), www.computerlanguage.com

Laplink Gold, SyncToy, and Viceversa software (see also Section 57.3.4.4) check files on the source and target machines and allow one to choose how to transfer changes:

- *Clone* the target machine using the source machine as the standard (i.e., make the target identical to the source, including deletions).
- Copy all new files from the source to the target without deleting any files on the target.
- Copy all new files from either machine to the other without deleting files.

Software solutions also can provide automatic copying of data onto separate internal or external disks. There are dozens of such utilities available, many of them free.[4] Examples include:

- SureSync software, which runs on most versions of Windows, and which can replicate and synchronize files on Windows, MAC, and Linux.
- NFTP 1.64, which runs on Windows, OS/2, BeOS, Linux, and UNIX; it includes "transfer resume, automatic reconnect, secure authentication, extensive firewall and proxy support, support for many server types, built-in File Transfer Protocol (FTP) search, mirroring, preserving timestamps and access rights, control connection history."[5]
- UnixWare Optional Services include Disk Mirroring software for what was originally the Santa Cruz Operation (SCO) UNIX.
- Double-Take software goes beyond periodic backup by capturing byte-level changes in real time and replicating them to an alternate server, either locally or across the globe.

Early users of software-based disk mirroring suffered from slower responses when updating their primary files because the system had to complete output operations to the mirror file before releasing the application for further activity. However, today's software uses part of system memory as a buffer to prevent performance degradation. Secondary (mirror) files are nonetheless rarely more than a few milliseconds behind the current status of the primary file.

57.2.4 Logging and Recovery. If real-time access to perfect data is not essential, a well-established approach to high-availability backups is to keep a log file of all changes to critical files. *Roll-forward* recovery requires:

- Backups that are synchronized with log files to provide an agreed-on starting point
- Markers in the log files to indicate completed sequences of operations (called *transactions*) that can be recovered
- Recovery software that can read the log files and redo all the changes to the data, leaving out incomplete transactions

An alternative to roll-forward recovery is *roll-backward* recovery, in which diagnostic software scans log files and identifies only the incomplete transactions and then returns the data files to a *consistent state*. For a detailed discussion of logging, see Chapter 53 in this *Handbook*.

Microsoft Windows XP and later versions include the *System Restore* function.[6] System Restore uses a disk buffer of at least 200 MB to store several time-stamped copies of critical executable files and settings, including system registry settings, when users install new software or patches. Users can also create restore points themselves. Users can thus return their operating system to a state consistent with that of a prior time. In addition, Windows offers a device driver rollback function for a single level of driver recovery.[7]

57.2.5 Backup Software. All operating systems have utilities for making backups. However, sometimes the utilities that are included with the installation sets are limited in functionality; for example, they may not provide the full flexibility required to produce backups on different kinds of removable media. Generally, manufacturers of removable media include specialized backup software suitable for use with their own products.

There are so many backup products on the market for so many operating environments that readers will be able to locate suitable candidates easily with elementary searches. When evaluating backup software, users will want to check for these minimum requirements:

- The software should allow complete control over which files are backed up. Users should be able to obtain a report on exactly which files were successfully backed up and detailed explanations of why certain files could not be backed up.
- Restore operations should be configurable to respect the read-only attribute on files or to override that attribute globally or selectively.
- Data compression should be available.
- A variety of standard, open (not proprietary) data encryption algorithms should be available.
- Backups should include full directory paths on demand, for all files.
- Scheduling should be easy, and not require human intervention once a backup is scheduled.
- Backups must be able to span multiple volumes of removable media; a backup must not be limited to the space available on a single volume.
- If free space is available, it should be possible to put more than one backup on a single volume.
- The backup software must be able to verify the readability of all backups as part of the backup process.
- It should not be easy to create backup volumes that have the same name.
- The *restore* function should allow selective retrieval of individual files and folders or directories.
- The destination of restored data should be controllable by the user.
- During the restore process, the user should be able to determine whether to overwrite files that are currently in place; the overwriting should be controllable both with file-by-file confirmation dialogs and globally, *without* further dialog.

57.2.6 Removable Media. The density of data storage on removable media has increased thousands of times in the last half century. For example, in the 1970s, an

8-inch word-processing diskette could store up to 128 KB. In contrast, at the time of writing (April 2008), a removable disk cartridge 3 inches in diameter stored 240 GB, provided data transfer rates of 30 MB per second, and cost around $100. Moore's Law predicts that computer equipment functional capabilities and capacity double every 12 to 18 months for a given cost; this relationship definitely applies to mass storage and backup media.[8]

57.2.6.1 External Hard Disk Drives. Many styles of portable, removable stand-alone disk drives with universal serial bus (USB) and FireWire connections are available at relatively low cost. For example, at the time of writing (July 2013), Western Digital USB 3.0 portable drives with 500 GB to 2 TB storage fit easily into a shirt pocket to provide backup and recovery for travelers with laptop computers and cost from $70 to $160.

For office or higher-volume applications, a 3 TB Western Digital MyBook USB 3.0 external hard drive cost about $130.

Such devices make it easy to create one or more encrypted clones (exact copies, although not necessarily bit-for-bit copies)[9] of all of one's laptop or tower computer data, and can allow one to synchronize the clones as often as necessary. The disadvantage compared with using multiple independent storage media is that if the unit fails, all the backup data are lost. However, with costs so low, it may be affordable to have several units and to synchronize a different one at every backup, allowing a collection of backups to accumulate.

57.2.6.2 Removable Hard Disk Drives. Hard disks configured so that they can slide into bays provide for both expansion of disk storage and for backups of critical data. For example, in July 2013, the Imation RDX Removable HDD Storage System USB 3.0 was selling for about $300 for a docking station plus one 750-GB cartridge; a 1.5 TB cartridge cost about $265. One of the advantages of these systems is that one can connect install bays for multiple drives, allowing concurrent I/O to multiple cartridges, thus increasing the speed of backup.

57.2.6.3 Optical Storage. Many users and systems now use optical storage for backups. Single-layer DVDs can store 4.7 GB, while dual-layer DVDs can store 8.5 GB on media that cost less than a dollar per disk. Write parameters in the optical drivers allow users to specify whether they intend to add to the data stored on disks or whether the disks should be closed to further modification at the end of the write cycle. Internal and external optical read/write drives cost as little as $20.

Blu-ray discs™ (BDs) are an optical disk alternative to the DVD format that won out over the competing High Definition (HD) DVD standard in 2008. Single-layer BDs can store 25 GB in the same size disk as CDs and DVDs; dual-layer BDs store 50 GB. BDXL triple-layer drives can store 100 GB per disk and BDXL quad-layer drives can store 128 GB per disk.

An upcoming technology to monitor is the *holographic* disk, with storage capacities of up to 1.6 TB. However, at the time of writing in July 2013, the technology was still viewed as experimental. The hVault company, a startup in 2011, was in beta-test of these devices for the commercial market (not personal use) in configurations supporting petabytes of storage. The company described how the system works:

> Holographic Storage is a totally new way of storing digital data. It works by splitting a blue laser beam in two, a reference beam and a data beam. The data beam is modulated with

a spatial light modulator (like in a digital projector) that creates a very high density array of 1.5 million pixels. Where the data beam and the reference beam intersect, a holographic interference pattern is formed. When focused on an unexposed, light sensitive holographic disk, a permanent 3D holographic image is formed in the light sensitive layer of the disk. To recover the data from the disk, a reference laser is focused on the disk and an image of the original digital array is formed and captured on a CCD image sensor.

Once the image containing 1.5Mbits of data is written to the disk, the angle of the reference beam is moved slightly and another 1.5Mbits can be written in the same 3D volume as the first image. This process is repeated many times with slightly different reference beam angles to create multiple, non-interfering data images in the same volume of the disk. Then the disk is rotated slightly to the next area of the disk and the multiple write sequence is repeated. Unlike magnetic Hard Drives and optical disks such as CD, DVD and Blu-Ray, where bits are written one bit at a time onto a spinning disk, holographic data storage reads and writes 1.5Mbits at a time on a stationary disk. This results in improved read and write speeds as well as much lower power consumption because there is no need to spin the disk constantly.

Once the disk is exposed with the data to be saved, the disk is cured with the laser which is analogous to developing and fixing photographic film. Once developed, black and white photographic film is no longer light sensitive and is relatively inert, which gives it very good long term image retention. Similarly, holographic disks become inert once they are cured they become unalterable by natural elements, human errors or tampering. Accelerated life testing indicates that holographic disks can store their digital data safely for over 50 years with no loss of data.[10]

hVault products were described as follows:

hVault provides complete holographic storage systems for any size archive. For small archives, the single drive, 15 disk slot, rack mount autoloader is a very cost effective way to create, store and retrieve Terabytes of digital data in compact, low cost unit. Larger archives can opt for the automated libraries that can be configured with 2 to 8 drives and up to 540 disk slots. Expansion cabinets can be added to increase the total slot count to 2140 disks. These robotic units provide very high system flexibility due to the ability to read and write simultaneously from multiple drives. hVault holographic storage systems provide expansion capability for your growing archive storage needs. hVault provides complete holographic storage systems for any size archive. For small archives, the single drive, 15 disk slot, rack mount autoloader is a very cost effective way to create, store and retrieve Terabytes of digital data in compact, low cost unit. Larger archives can opt for the automated libraries that can be configured with 2 to 8 drives and up to 540 disk slots. Expansion cabinets can be added to increase the total slot count to 2140 disks. These robotic units provide very high system flexibility due to the ability to read and write simultaneously from multiple drives. hVault holographic storage systems provide expansion capability for your growing archive storage needs.

All hVault storage systems are designed to plug in to existing network architectures with no disruption to existing workflows. Flexible control and data path options make integration easy. Whether your network uses Gigabit Ethernet or FibreChannel, hVault libraries are plug and play. Flexible control interfaces to the robotic controls make integration into your control architecture a snap.

With any hVault storage system, ANY content on ANY disk is accessible in less than 10 seconds. Once the disk is loaded into a drive, the read or write throughput is 20MB/sec. Caching provides a very high bandwidth input channel that allows bursting of very high speed data to the system.[11]

Prices were projected to start at about $50,000 for the low-end 15-disk autoloader. There were indications that the company might also offer consumer units capable of storing more than 1 TB per platter.

57.2.6.4 Tape Cartridge Systems. The old 9-track, reel-to-reel 6,250 bytes-per-inch (bpi) systems used in the 1970s and 1980s held several hundred MB. Today's pocket-size tape cartridges hold GBs. For example, the industry leader in this field, a

Sony LTO3 cartridge, has 400-GB native and 800-GB compressed maximum storage capacity per cartridge and a transport speed of 80 MB/s native and 160 MB/s compressed data transfer. Cartridges have mean-time-between-failure (MTBF) of 250,000 hours with 100 percent duty cycles and can tolerate 1 million tape passes. All such systems have streaming input/output (I/O) using random access memory (RAM) buffers to prevent interruption of the read/write operations from and to the tapes and thus to keep the tape moving smoothly to maximize data transfer rates.

Although this buffering is effective for directly attached tape systems, this is not effective for backing up data over an LAN. When backing up data over an LAN, users will probably experience *shoeshining* of the tape drive (rapid back-and-forth movement of the tape), which reduces the effective throughput rate significantly. For example, a 50 MB per second (ps) tape drive, receiving a data stream at 25 MBps, does not write at 25 MBps. Rather it fills a buffer and writes short bursts at 50 MBps. The mechanical motion of the tape moves empty tape over the read/write head after the data transfer ends. The tape motor then stops, rewinds to reposition the tape at the appropriate place on the tape for continued writing, and prepares itself to write another short burst at 50 MBps. This repositioning process of stopping, rewinding, and getting back up to speed is called *backhitching*. Each backhitch can take as long as a few seconds. Frequent backhitching is called *shoeshining* because the tape activity mimics the movement of a cloth being used to shine shoes. The farther one operates from the designed throughput rate, the less time one spends writing data and the more time backhitching. A 50-MBps tape drive that is receiving data at 40 MBps will actually write at 35 MBps because it is spending at least 20 percent of its time shoeshining. Similarly, a data stream at 30 MBps will actually be written at 20 MBps, a 20 MBps data stream at 10 MBps, and a 10 MBps data stream at less than 1 MBps.

In conjunction with automated tape library systems, holding many cartridges and capable of switching automatically to the next cartridge, tape cartridge systems are ideal for backing up servers and mainframes with TBs of data. Small library systems keep 10 to 20 cartridges in position for immediate access, taking approximately 9 seconds for an exchange. These libraries have approximately 2 million mean exchanges between failures, with MTBF of around 360,000 hours at 100 percent duty cycle.

The large enterprise-class library systems such as the 8-frame SpectraLogic's T950 can be configured with up to 10,050 slots and 120 LTO (Linear Tape-Open, the open-software alternative to DLT or Digital Linear Tape) drives. The 8-frame T950 offered a native throughput of 14.4 GB/second using LTO-4 drives. These library systems can encrypt, compress, and decrypt data using the strongest version of the Advanced Encryption Standard (AES), to meet regulatory and security compliance requirements. As another example, the Plasmon Ultra Density Optical (UDO) Library offered up to 19 TB of archival storage capacity for enterprise systems at a price of about $100,000.

57.2.6.5 Virtual Tape Library. A virtual tape library (VTL) is a disk-based system that emulates a tape library, tape drives, and media. These systems are used in disk-based backup and restore operations.[12] VTLs used for backup emulate tape drive and library functions and can provide a transition for legacy systems to an integrated disk-based backup and recovery. These units (e.g., the IBM Virtual Tape Server introduced in 1997) provide backward compatibility for older production systems by using optical, or modern high-capacity magnetic disk storage but presenting the data as if they were stored on magnetic tapes and cartridges. VTLs can use existing storage architecture and can be used remotely. Naturally, their throughput is far greater than the older tape-based backup systems.

As an example of technology available at the time of writing (July 2013), FalconStor VTL was described by the manufacturer as follows:

Ultra-fast, reliable backup and deduplication

There is a reason why FalconStor® Virtual Tape Library (VTL) is the industry's market-leading virtual tape solution: It provides enterprise-level disk-based backup and deduplication to organizations of all sizes at the fastest available speeds. FalconStor VTL optimizes backup speed, performance, and reliability without interrupting existing backup environments, minimizing backup windows and making data more available and secure.

Optimized

- Flexible, policy-based deduplication: Inline, post-process, concurrent, or no deduplication
- Fastest backup speeds: Up to 5.8TB/hr
- Fastest deduplication speeds: Up to 40TB/hr

Available

- Enhanced replication ensures efficient bandwidth usage in LAN and WAN environments
- Global deduplication scales to over 2PB of usable storage
- High-availability (HA) cluster configurations for large enterprises

Secure

- Sophisticated physical tape integration; automated, nondisruptive direct tape export
- Replication with Advanced Encryption Standard (AES) encryption
- Federal Information Processing Standards (FIPS) compliance support

Optimized, high-performance disk backup

By offering ultra-fast backup of up to 1.6GB/sec (5.8TB/hr), FalconStor VTL addresses the single biggest backup challenge: the backup window. When eight appliances are combined into one logical unit, speeds accelerate to 12.8GB/sec. Combined with high-speed protocols such as 8Gb Fibre Channel (FC) and 10Gb Ethernet, FalconStor VTL can sustain deduplication speeds of up to 40TB/hr—the fastest in the industry.

Highly flexible, FalconStor VTL enables you to set customized policies for deduplication, replication, and archiving. You can select from up to four deduplication types—inline, post-process, concurrent, or no deduplication—on a per-job basis.

Global deduplication ensures availability with efficiency

FalconStor VTL ensures the availability of critical data. It replicates data from across the globe to a central clustered repository of unique, deduplicated data. WAN-optimized replication ensures fast and reliable backup and recovery while reducing bandwidth requirements. Global deduplication simplifies storage management and reduces capacity requirements by up to 95%, shrinking the storage footprint at the data center.

Secure backup, recovery, and archiving

FalconStor VTL provides support for a broad range of systems, tapes, and tape libraries, including Symantec OST environments. FalconStor VTL ensures the safety and security of critical data through features such as extended tape management, encryption with Advanced

Encryption Standard (AES), tape shredding, and Federal Information Processing Standards (FIPS) compliance support.

Flexible deployment options

FalconStor VTL is available in a variety of form factors, providing enterprise-class functionality to customers ranging from small-to-medium businesses (SMB) to large global enterprises with multiple remote and branch offices (ROBO)… . [13]

VTLs may be offered by certain vendors as software only, purpose-built appliances, or virtual appliances for VMware environments. Pricing for virtual appliances generally start at $2,000–$3,000 per license and an additional cost for capacity.

57.2.6.6 Personal Storage Devices: Flash Memory. USB flash drives have become ubiquitous due to the decline in price and the vast amount of data that they can store. Cheap multi-MB and multi-GB USB flash drives (also called *memory sticks* and *thumb drives*) are available at drugstores, are given away at conferences, and are available in a variety of forms, such as pens, pocketknives, and supposedly amusing variants such as dolls and even sushi. Other personal devices using flash memory can also be used to store large amounts of data. These include music players, personal digital assistants, and mobile phones. For example, Secure Digital (SD) memory cards with 256 GB of storage were available at the time of writing (July 2013) for about $600; a 128-GB card sold for around $300; and a 64-GB card cost about $100.

Although these devices can provide a convenience for the user who wants to transfer large amounts of data, they pose a serious security risk to businesses, since sensitive company data can be easily removed, against company policy. Although some USB flash drives have built-in encryption to protect the contents from being exposed if the device is lost or stolen, most flash drives and other personal devices store data in unencrypted form—a serious problem if the device is lost or stolen. Some companies have implemented technical means (*data-loss prevention*, or DLP) to block the use of USB flash and other devices from downloading data from corporate PCs. Just as these devices can download data from a computer, they can also be a vector for transferring malware to a computer and to a network. User policies and education are the primary defenses to supplement technical controls.

Flash drives are sometimes recommended as a backup medium for end users. However, they suffer from a number of disadvantages:

- They are usually physically small and therefore easy to misplace or lose
- They are relatively fragile (e.g., their enclosures break easily if struck and may not be water resistant)
- They are relatively expensive per storage unit compared with magnetic or optical media

57.2.6.7 Developing Storage Technologies. IBM research scientists are working on *Millipede*, a method of storing data using an atomic-force microscope to punch nanometer-sized depressions in thin polymer films.[14] The storage density is about 100 GB per square cm.[15] It will be interesting to see if this technology develops into yet another order-of-magnitude increase in storage capacity and precipitous decline in storage costs.

Another astounding leap forward is the demonstration of multilayered nanostructured quartz glass, which has allowed storage of 360TB on a glass disk with expected lifetimes so long they have not yet been estimated accurately.[16]

57.2.7 Labeling. Regardless of the size of a backup, every storage device, from diskettes to tape cartridges to external or removal disk drives, should be clearly and unambiguously tagged, both electronically and with adhesive labels. Backups to hard drives can be placed in folders or files with as much detail as required; for example, *BU_2016-03-31_FULL* or *BU_2016-04-11_DB.ZIPX*.

Larger-capacity media, such as cartridges used for UNIX and mainframe systems, often have extensive bar-code or electronic labeling available. On some systems, it is possible to request specific storage media and have the system automatically refuse the wrong media if they are mounted in error. Tape library systems typically use optical bar codes that are generated automatically by the backup software and then manually or automatically affixed to each cartridge for unique identification. Magnetic tapes and cartridges have electronic labels written onto the start of the recording medium in each unit, with specifics that are particular to the operating system and tape-handling software.

An unlabeled storage medium or one with a flimsily attached label is evidence of a bad practice that will lead to confusion and error. Sticky notes, for example, are not a good way to label diskettes and removable disks. If the notes are taken off, they can get lost; if they are left on, they can jam the disk drives. There are many types of labels for storage media, including printable sheets suitable for laser or inkjet printers and using adhesive that allows removal of the labels without leaving a sticky residue. At the very least, an exterior label should include this information:

- Date the volume was created (e.g., *2016–09–08*)
- Originator (e.g., *Bob R. Jones, Accounting Dept.*)
- Description of the contents (e.g., *Engineering Accounting Data for 2016*)
- Application program (e.g., *FastAccounting v3.4.5*)
- Operating system (e.g., *Windows 9 v3.4.5.6*)

Storing files with *canonical names* (names with a fixed name or a defined structure) on the media themselves is also useful. An example of canonical files much used in installation packages is *READ.ME* or *README.TXT*. For example, an organization can mandate files such as the following as minimum standards for identifying and documenting its storage media:

- *ORIGIN.mmm* (where *mmm* represents a sequence number for unique identification of the storage set) indicating the originating system (e.g., the *ORIGIN.1234* file could contain data such as *Accounting Workstation number 3875-3* or *Bob Whitmore's SPARC in Engineering Rm 379 serial 98765*). This file can be written to every volume.
- *DATE.mmm* showing the date (preferably in year-month-day sequence) and time (usually 24-hour clock) on which the storage volume was created (e.g., *DATE.1234* could contain the data *2016–09–08 23:41*). Such a file can be written to each volume of the storage set.

- *SET.mmm* to describe exactly which volumes are part of a particular set; e.g., SET.1234 could include *SET 1234: VOL 444, VOL 445, VOL 446.* This kind of file is normally located on the last volume of the set.
- *INDEX.mmm*, an index file on the last volume listing all the files on all the volumes of that particular storage set; for example, INDEX.1234 could contain *SET 1234: VOL 444 – FIL F1, F2, F3; VOL 445 – FIL F4, F5, F6, F6; VOL 446 – FIL F7, F8.*
- *VOLUME.nnn* (where *nnn* represents a sequence number for unique identification of the medium) that contains an explanation; e.g., *VOLUME.444* could include data such as *VOL 444 SET 1232 NUMBER 1.*
- *FILES.nnn*, which lists all the files on that particular volume of that particular storage set; for example, contents of file FILES.444 could include *SET 1234: VOL 444 – Files F1, F2, F3.*

Such labeling is best handled by an application program; many backup programs automatically generate analogous files.

57.2.8 Indexing and Archives. As implied earlier, backup volumes need a mechanism for identifying the data stored on each medium. Equally important is the capacity to locate the storage media where particular files are stored; otherwise, one would have to search serially through multiple media to locate specific data. Although not all backup products for personal computers (PCs) include such functionality, many do; server and mainframe utilities routinely include automatic indexing and retrieval. These systems allow the user to specify file names and dates, using wild-card characters in the defined identification pattern to signify ranges and also to display a menu of options from which the user can select the appropriate files for recovery.

This file-level indexing, though, is not sufficient for retrieving unstructured content such as word processing documents, e-mails, instant messages, diagrams, Web pages, and images that may be required for regulators, business research, and legal subpoenas (e-discovery). The Federal Rules of Civil Procedure (FRCP) specifically address discovery and the duty to disclose evidence in preparation for trial.[17] Recent changes to FRCP specifically address electronic data.

Due to the cost of locating and retrieving data for litigation, companies are changing the way they manage the life cycle of data. For instance, the cost to do a forensics search of one backup tape to locate e-mails that contain certain relevant content can be $35,000 per tape. In some recent cases, a search of e-mail tapes has involved over 100 tapes. Companies are now implementing archiving systems to address this requirement.

A method of providing fast access to fixed content (data that are not expected to be updated) is *content-addressed storage* (CAS). CAS assigns content a permanent place on a disk. CAS stores content so that an object cannot be duplicated or modified once it has been stored; thus, its location is unambiguous.

When an object is stored in CAS, the object is given a unique name that also specifies the storage location. This type of address is called a *content address*. It eliminates the need for a centralized index, so it is not necessary to track the location of stored data. Once an object has been stored, it cannot be deleted until the specified retention period has expired. In CAS, data is stored on disk, not tape. This method streamlines the process of searching for stored objects. A backup copy of every object is stored to enhance reliability and to minimize the risk of catastrophic data loss. A remote

monitoring system will notify the system administrator in the event of a hardware failure.

Combined archiving and backup systems can reduce the amount of stored data by reducing the amount of duplicate data, thereby reducing the storage cost and time to retrieve or restore data files.

EMC's Centera product first released in 2002 is an example of a CAS product. The product overview highlights the following features:

- **Content authenticity**: Centera ensures that your online, long-term archive data is authentic, cannot be overwritten, and is immutable with write-once, read-many (WORM) technology.
- **Governance and compliance**: Meet industry regulations and corporate requirements using a wide range of retention capabilities on a hardened platform with assured content authenticity and legal admissibility of content.
- **High availability**: Ensure data accessibility with redundant array of independent nodes (RAIN), flexible protection schemes, self-healing, and replication for unmatched protection and 99.999 percent availability.
- **Efficient operation**: Reduce total cost of ownership (TCO) with seamless, nondisruptive scale and location-independent storage. Automate storage management tasks including self-configuring, self-healing, and virtual storage pools. Ensure low-capacity overhead with single-instance storage.
- **Broad partner ecosystem**: Extend Centera investment value with more than 300 out-of-the box, integrated archiving applications to manage email, file, medical imaging, content management, video, and voice archiving on a single archiving platform.[18]

57.2.9 Registry Backups. On Windows systems, the central repository of settings for all installed programs and device drivers is the *registry*. The equivalent for Mac computers is the *system preferences*. On *nix systems, similar information is in files in specific folders such as */settings*.

For some operating systems, it may not be possible to make a usable copy of such data using file-system–based backup software. However, there may be options in some backup software to make bit-for-bit copies of the registry (or equivalent) so that settings can be restored to a system that has suffered catastrophic damage or that has been restored to a previous condition. Under Windows, a variety of registry-editors typically include options for automatically making restorable copies of the registry. Windows itself has functions for so doing; for example, the *Create a restore point* function allows one to store usable copies of the registry. The *Restore system files and settings from a restore point* or *Restore your computer to an earlier time* are pointers to returning to what one hopes will be a better version of the system configuration.

When creating restore point, it is wise to take advantage of options for storing information explaining why the restore point has been taken; for example, a description could be *2016-06-02 before installing BaffleGab 4.3.6.*

57.2.10 Bootable Versions. Registry backups can be useful, but what does one do if the operating system disk has been completely corrupted or if one has to initialize a new computer system? On Windows systems with hundreds of installed programs, recreating a work environment can take days of annoying installations,

re-registration, and repeated boot cycles. To avoid such time-wasting labor, one should keep bootable versions of the entire system.

Most system-backup software today includes the option to allow restoration of the entire configuration, including the registry, every installed program, and all installed options for settings of those programs. Some packages allow access to individual folders and files for restoration and others don't; however, the key is that a bootable backup can save days of time and effort.[19]

The ultimate bootable backup is the *image backup* which is a bit-for-bit copy of the disk containing operating system files. Such files are commonly used in corporate environments to ensure easy restoration of all networked workstations to an authorized initial state. They can save enormous amounts of time for large networks by eliminating the need to install upgrades manually system by system; for example, a university network with 1,000 lab computers accessible to students could not possibly be managed without image backups to centralize the configuration updates of all those machines.

57.2.11 Cloud Backup. An alternative to making onsite backup copies is to pay a third party to make automatic backups via high-speed telecommunications channels and to store the data in a secure facility. Some firms involved in these services move data to magnetic or optical backup volumes, but others use RAID (see Section 57.2.3.1) for instant access to the latest backups. Additional features to look for when evaluating online backup facilities:

- Compatibility of backup software with the computing platform, operating system, and application programs
- Availability of different backup options: full, differential, incremental, and delta
- Handling of files that are held in an *open* state by application programs
- Availability and costs of sufficient bandwidth to support desired data backup rates
- Encryption of data during transmission and when stored at the service facility
- Strong access controls to limit access to stored data to authorized personnel
- Physical security at the storage site, and other criteria similar to those listed in Section 57.5.3.5
- Methods for restoring files from these backups, and the speed with which they can be accomplished

In a May 2013 forecast on global network bandwidth growth, Cisco's "Visual Networking Index" included the following predictions, among many others:

- Annual global IP traffic will surpass the zettabyte threshold (1.4 zettabytes) by the end of 2017. In 2017, global IP traffic will reach 1.4 zettabytes per year, or 120.6 exabytes per month. Global IP traffic will reach 1.0 zettabytes per year or 83.8 exabytes per month in 2015.
- Global IP traffic has increased more than fourfold in the past 5 years and will increase threefold over the next 5 years. Overall, IP traffic will grow at a compound annual growth rate (CAGR) of 23 percent from 2012 to 2017.
- Nearly half of all IP traffic will originate with non-PC devices by 2017. In 2012, only 26 percent of consumer IP traffic originated with non-PC devices, but by 2017 the non-PC share of consumer IP traffic will grow to 49 percent. PC-originated

traffic will grow at a CAGR of 14 percent, while TVs, tablets, mobile phones, and machine-to-machine (M2M) modules will have traffic growth rates of 24 percent, 104 percent, 79 percent, and 82 percent, respectively.

- Traffic from wireless and mobile devices will exceed traffic from wired devices by 2016. By 2017, wired devices will account for 45 percent of IP traffic, while Wi-Fi and mobile devices will account for 55 percent of IP traffic. In 2012, wired devices accounted for the majority of IP traffic at 59 percent.[20]

The global bandwidth of the Internet has been climbing steadily and is predicted to continue accelerated growth over the remainder of the 2010 decade:

> …[I]nternational bandwidth availability has soared ("used bandwidth" refers to the capacity deployed by providers, rather than bandwidth consumed by end users). From 1.4 terabits per second in 2002, it steadily climbed to 6.7 terabits in 2006 and has now reached 92.1 terabits per second. TeleGeography expects that number to hit 606.6 terabits per second in 2018 and 1,103.3 terabits per second in 2020.[21]

The increasing availability of higher bandwidth connections is encouraging the growth of cloud-based file-sharing and backup services, especially for mobile computing platforms (laptops, tablets, phones). Although some services date back for decades (e.g., Asigra[22]), dozens of companies and their remote-backup services have appeared in the twenty-first century to date. The Storage Networking Industry Association (SNIA) convened a Cloud Storage Initiative (CSI) in 2010; members included Cisco Systems, Cleversafe, Inc., EMC, Evaluator Group, Hewlett-Packard, Hitachi Data Systems, IBM, Inktank, Mezeo Software, NetApp, Oracle, Scality, SpectraLogic, SwiftTest, Symantec, and Vedams, as of July 3, 2013.[23]

The SNIA CSI has defined the *Cloud Data Management Interface* (CDMI) cloud-storage standard, which

> …defines the functional interface that applications will use to create, retrieve, update, and delete data elements from the Cloud. As part of this interface the client will be able to discover the capabilities of the cloud storage offering and use this interface to manage containers and the data that is placed in them. In addition, metadata can be set on containers and their contained data elements through this interface.
>
> This interface is also used by administrative and management applications to manage containers, accounts, security access, and monitoring/billing information, even for storage that is accessible by other protocols. The capabilities of the underlying storage and data services are exposed so that clients can understand the offering.[24]

Virtualization has also been implemented in cloud backups. For example, in early 2012,

> Amazon Web Services…launched a public beta test of AWS Storage Gateway, which allows enterprises to back up application data in Amazon's cloud using a software appliance…. The Storage Gateway appliance is a virtual machine that runs on VMware's virtualization software. It uses an iSCSI interface to integrate with applications. The appliance stores data on local storage hardware, while uploading backup snapshots to Amazon's cloud. This provides low-latency access to data and off-site backups in the cloud….
>
> Enterprises that want to use Storage Gateway pay a monthly fee of US$125 per activated gateway. Added to that is the cost of storage capacity used by the snapshots, and traffic out of Amazon's cloud.
>
> The snapshots are stored and billed as Amazon EBS (Elastic Block Store) snapshots, which cost from $0.14 per GB per month.

> The cost of data transfer drops as the amount of data increases: Up to 10 TB of transferred data in one month costs $0.12 per GB, and the next 40 TB costs $0.09 per GB. Rate tiers are chosen based on aggregate usage across multiple Amazon services, including EC2 (Elastic Compute Cloud), S3, and VPC (Virtual Private Cloud).[25]

One of the critical issues when choosing a cloud-backup service is quality of service: If the service goes offline, there can be serious problems both updating the backups and accessing backup files.[26]

Another consideration when deciding whether to use cloud-backup services is the consequences for court-ordered forensic seizure of the data. Gordon Merrill published a list of questions that should be discussed with possible vendors:

- Do we even know which specific drives were in use by XYZ before the crash at EC2?
- Would Amazon have the ability to remove those drives and replace with others if ordered to do so?
- How many other companies' data have been written on those drives in the interim?
- If the original XYZ data have been overwritten by other companies and the drives are removed for recovery attempts, does the removal mean that the later users have now lost control of their data?
- Do the current users of the removed drives have to be served with a notice that the drives are being forensically reviewed?
- Is there a legal requirement that the current users need to be notified?
- Are the current users due a description of how their data was handled during the recovery and how it was destroyed when the exam was complete in order for them to produce the same to their customers as ordered for compliance with applicable laws?

One last concern facing most companies legally is that of legal hold orders and/or search warrants.

- If XYZ is being investigated by the Department of Justice (DoJ) and they want to find out more during an investigation, can the DoJ serve a warrant to Amazon and search without ever notifying XYZ that the search is going on?
- If the same hard drives are now in use by company ABC, does ABC get notified of the search and seizure or is the warrant on Amazon enough to search without any notice to the companies involved?[27]

Finally, shared folders such as those on Dropbox must be backed up regularly in a separate way—whether on the local hard drives or using a different cloud-backup service. If any user of the shared folders deletes files, all the users will have their copies automatically deleted. Without explicit backups, valuable information can be destroyed.[28]

For an extensive review of cloud storage and backup, see the paper by Legendre.[29]

57.3 BACKUP STRATEGIES. There are different approaches to backing up data. This section looks at what kinds of data can be backed up and then reviews appropriate

ways of choosing and managing backups for different kinds of computer systems. Some of the challenges that a backup strategy should address are:

- Rapid growth of electronic data
- Shrinking time for backing up data
- Increasing backup costs
- Performance bottlenecks
- Compliance and legal discovery

57.3.1 Exclusive Access. All backup systems have trouble with files that are currently in use by processes that have opened them with write access (i.e., which may be adding or changing data within the files). The danger in copying such files is that they may be in an *inconsistent state* when the backup software copies their data. For example, a multiphase transaction may have updated some records in a detail file, but the corresponding master records may not yet have been posted to disk. Copying the data before the transaction completes will store a corrupt version of the files and lead to problems when they are later restored to disk.

Backup software usually generates a list of everything backed up and of all the files *not* backed up; for the latter, there is usually an explanation or a code showing the reason for the failure. Operators always must verify that all required files have been backed up and must take corrective action if files have been omitted.

Workstation backups for operating systems such as Windows and Mac usually include options for using what is often called *volume shadow copy services* to provide access to files in current use. The Windows Volume Shadow Copy Service (VSS) is a feature of Windows 7, Windows Server 2003, Windows XP, and Windows Vista. VSS was removed from Windows in the development of Windows 8.[30]

Some high-speed, high-capacity backup software packages provide a buffer mechanism to allow high-availability systems to continue processing while backups are in progress. In these systems, files are frozen in a consistent state so that backup can proceed, and all changes are stored in buffers on disk for later entry into the production databases. However, even this approach cannot obviate the need for a minimum period of quiescence so that the databases can reach a consistent state. In addition, it is impossible for full functionality to continue if changes are being held back from the databases until a backup is complete; all *dependent* transactions (those depending on the previously changed values of records) also must be held up until the files are unlocked.

57.3.2 Types of Backups. Backups can include different amounts and kinds of data:

- *Full* backups store a copy of everything that resides on the mass storage of a specific system. To restore a group of files from a full backup, the operator mounts the appropriate volume of the backup set and restores the files in a single operation.
- *Differential* backups store all the data that have changed since a specific date or event; typically, a differential backup stores everything that has changed since the last full backup. The number of volumes of differential backups can increase with each additional backup. To restore a group of files from a differential backup, the

operator needs to locate the latest differential set and also the full backup on which it is based to ensure that all files are restored. For example, suppose a full backup on Sunday contains copies of files *A, B, C, D*, and *E*. On Monday, suppose that files *A* and *B* are changed to *A′* and *B′* during the day; then the Monday evening differential backup would include only copies of files *A′* and *B′*. On Tuesday, suppose that files *C* and *D* are changed to *C′* and *D′*; then the Tuesday evening backup would include copies of files *A′, B′, C′*, and *D′*. If there were a crash on Wednesday morning, the system could be rebuilt from Sunday's full backup (files *A, B, C, D,* and *E*), and then the Tuesday differential backup would overwrite the changed files and restore *A′, B′, C′*, and *D′*, resulting in the correct combination of unchanged and changed files (*A′, B′, C′, D′*, and *E*).

- *Incremental* backups are a more limited type of differential backups that typically store everything that has changed since the previous full or incremental backup. As long as multiple backup sets can be put on a single volume, the incremental backup requires fewer volumes than a normal differential backup for a given period. To restore a set of files from incremental backups, the operator may have to mount volumes from all the incremental sets plus the full backup upon which they are based. For example, using the same scenario as in the explanation of differential backups, suppose a full backup on Sunday contains copies of files *A, B, C, D,* and *E*. On Monday, suppose that files *A* and *B* are changed to *A′* and *B′* during the day; then the Monday evening incremental backup would include only copies of files *A′* and *B′*. On Tuesday, suppose that files *C* and *D* are changed to *C′* and *D′*; then the Tuesday evening backup would include copies only of files *C′* and *D′*. If there were a crash on Wednesday morning, the system could be rebuilt from Sunday's full backup (files *A, B, C, D,* and *E*) and then the Monday incremental backup would allow restoration of *A′* and *B′* and the Tuesday incremental backup would restore *C′* and *D′*, resulting in the correct combination of unchanged and changed files (*A′, B′, C′, D′*, and *E*).
- *Delta* backups store only the portions of files that have been modified since the last full or delta backup; delta backups are a rarely used type, more akin to logging than to normal backups. Delta backups use the fewest backup volumes of all the methods listed; however, to restore data using delta backups, the operator must use special-purpose application programs and mount volumes from all the delta sets plus the full backup upon which they are based. In the scenario used earlier, the delta backup for Monday night would have records for the changes in files *A* in *delta-A* and records for the changes in file *B* stored in file *delta-B*. The Tuesday night delta backup would have records for the changes in files *C* and *D* stored in files *delta-C* and *delta-D*, respectively. To restore the correct conditions as of Tuesday night, the operator would restore the *A* and *B* files from the Sunday night full backup, then run a special recovery program to install the changes from the *delta-A* and a (possibly different) recovery program to install the changes from *delta-B*. Then it would be necessary to load the Tuesday night delta backup and to run recovery programs using files *C* and *D* with corresponding files *delta-C* and *delta-D*, respectively.
- *Versions* of working files are essential to protect the creators or users in case data are accidentally deleted without notice. For example, suppose an author constantly opens and saves all changes in a new document to filename *manuscript.docx* but uses no version numbers. At one point, the author accidentally deletes a significant block of text written the day before without noticing or under the belief that the

> deleted material is valueless. The author then continues working and eventually makes a backup—say, at the end of the week. There is therefore no backup of the deleted material: the *manuscript.docx* file overwrites the previous day's work. In contrast, the author may open yesterday's *manuscript_v32.docx* and immediately save it as *manuscript_v33.docx* for today's work. Even if a block of text written yesterday is deleted, it will be possible to open *manuscript_v32.docx* to recover the deleted text. In addition, having a consistent series of numbered, dated versions of a manuscript may serve to exonerate an author charged with plagiarism.[31]

Another aspect of backups is whether they include all the data on a system or only the data particular to specific application programs or groups of users:

- *System* backups copy everything on a system.
- *Application* backups copy the data needed to restore operations for particular software systems.

In addition to these terms, operators and users often refer to *daily, partial, weekly*, and *monthly* backups. These terms are ambiguous (they refer to the *frequency* of backups, not to the type) and could be used for any kind of backup; they should be used only in conjunction with the backup type (e.g., *monthly full backup* or *daily incremental backup*) when setting up procedures.

57.3.3 Computer Systems. Systems with different characteristics and purposes can require different backup strategies. This section looks at large production systems (mainframes), smaller computers used for distributed processing (servers), individual computers used primarily by one user (workstations), and portable computers (laptops) smartphones, or handheld computers and personal digital assistants (PDAs) which are still used by some corporations.

57.3.3.1 Mainframes. Large production systems using mainframes, or networks of servers, routinely do full system backups every day because of the importance of rapid recovery in case of data loss (see Chapters 42 and 43 of this *Handbook*). Using high-capacity tape libraries with multiple drives and immediate access to tape cartridges, these systems are capable of data throughput of up to 2 TB per hour. Some software vendors claim that they have been able to accelerate backups to 4 TB per hour. Symantec describes the increased throughput using their NetBackup Accelerator software as follows:

> One of the most painful parts of backup has always been the time required. It's a big deal—a recent Symantec survey revealed that 72 percent of enterprises would switch backup products if it would double their backup speed. Not content with simply doubling the speed, we implemented NetBackup Accelerator, which accelerates full backups by up to 100 times. This gives you the benefits of a full recovery image, instantly available, at the speed of performing an incremental backup.
>
> So what is the secret sauce? NetBackup Accelerator provides packaged in a single backup type the capability to read only changed files, dedupe these changed files, send only unique data from the client to the media server and from these unique changes, instantly make a full backup image available for recovery using NetBackup's synthesis engine.
>
> The results are astounding. Using old school backup, 1 TB of data could take up to 6 hrs to backup. With client deduplication alone, backup time is cut to 2 hours. With NetBackup Accelerator, backup time is cut to less than 15 minutes.[32]

Typically, all backups are performed automatically during the period of lowest system utilization. Because of the problems caused by concurrent access, mainframe operations usually reserve a time every day during which users are not permitted to access production applications. A typical approach sends a series of real-time messages to all open sessions announcing "Full Backup in xx minutes; please log off now." Operations staff have been known to phone offending users who are still logged on to the network when backups are supposed to start. To prevent unattended sessions from interfering with backups (as well as to reduce risks from unauthorized use of open sessions), most systems configure a time-out after a certain period of inactivity (typically 10 minutes). If users have left their sessions online despite the automatic logoff, mechanisms such as forced logoffs can be implemented to prevent user processes from continuing to hold production files open.

In addition to system backups, mainframe operations may be instructed to take more frequent backups of high-utilization application systems. Mission-critical transaction-processing systems, for example, may have several incremental or delta backups performed throughout the day. Transaction log files may be considered so important that they are also copied to backup media as soon as the files are closed. Typically, a log file is closed when it reaches its maximum size, and a new log file is initiated for the application programs.

57.3.3.2 Servers. Managers of networks with many servers have the same options as mainframe operations staff, but they also have increased flexibility because of the decentralized, distributed nature of the computing environment. Many network architectures allocate specific application systems or groups of users to specific servers; therefore, it is easy to schedule system backups at times convenient for the various groups. In addition to flexible system backups, the distributed aspect of such networks facilitates application backups.

Web servers, in particular, are candidates for high-availability clones of the entire Website. Some products provide for regeneration of a damaged Website from a read-only clone, one that includes hash totals allowing automatic integrity checking of all elements of the exposed and potentially damaged data on the public Website.[33]

57.3.3.3 Workstations. Individual workstations pose special challenges for backup. Although software and backup media are readily available for all operating systems, the human factor interferes with reliable backup. Users typically are not focused on their computing infrastructure; taking care of backups is not a high priority for busy professionals. Even technically trained users, who are aware of the dangers, sometimes skip their daily backups; many novice or technically unskilled workers do not even understand the concept of backups.

If the workstations are connected to a network, then automated, centralized, backup software utilities can protect all the users' files. However, with user disk drives containing as many as 500 GB of storage, and with the popularity of large files such as pictures and videos, storing the new data, let alone the full system, for hundreds of workstations can consume TBs of backup media and saturate limited bandwidths. It takes a minimum of 45 hours to transfer 1 TB over an OC-1 communications channel running at 51.84 megabits per second (Mbps); an OC48 running at 2,488.32 Mbps would transfer 1 TB in 56 minutes.[34] There are also privacy issues in such centralized backup if users fail to encrypt their hard disk files. In addition, it is impossible to compress encrypted data, so encryption and backup software should always compress *before* encrypting data.

57.3.3.4 Laptop Computers. A laptop computer is sometimes the only computer a user owns or is assigned; in other cases, the portable computer is an adjunct to a desktop computer. Laptop computers that are the primary system must be treated like workstations. Laptops that are used as adjuncts—for example, when traveling—can be backed up separately, or they can be synchronized with the corresponding desktop system.

Synchronization software offer a number of options to meet user needs:

- A variety of hardwired connection methods, including cables between serial ports, parallel ports, SCSI ports, and USB ports.
- Remote access protocols allowing users to reach their computer workstations via modem, or through TCP/IP connections via the Internet, to ensure synchronization or file access.
- Cloning, which duplicates the selected file structure of a source computer onto the target computer; cloning deletes files from the target that are not found on the source.
- Filtering, which prevents specific files, or types of files, from being transferred between computers.
- Synchronization, in which all changes on the source computer(s) are replicated onto the target computer(s). One-way synchronization updates the target only; two-way synchronization makes changes to both the target and the source computers.
- Compression and decompression routines to increase throughput during transfers and synchronizations.
- Data comparison functions to update only those portions of files that are different on source and target; for large files, this feature raises effective throughput by orders of magnitude.
- Security provisions to prevent unauthorized remote access to users' computers.
- Log files to record events during file transfers and synchronization.

In addition to making it easier to leave the office with all the right files on one's hard disk, synchronization of portable computers has the additional benefit of creating a backup of the source computer's files. For more complete assurance, the desktop system may be backed up daily onto one or two removable drives, at the same time that the portable computer's files are synchronized. If the portable is worked on overnight, all files should be synchronized again in the morning.

57.3.3.5 Mobile Computing Devices. Another area that is often overlooked is mobile computing devices: dedicated personal digital assistants (PDAs) and today's smart phones with calendars, contact lists, and so on. These devices often contain critically important information for their users, but not everyone realizes the value of making regular backups.

Mobile devices can usually be synchronized with a workstation file such as Outlook or Entourage. Synchronization programs for managing addresses, calendars, tasks, and so on can be configured to copy data unidirectionally from the workstation to the phone or from the phone to the workstation; bidirectional synchronization is also available. Android, Blackberry, iOS Nokia, Symbian, and Windows phones have a number of

widely used synchronizing apps available free or at modest cost through the appropriate app stores.

Synchronizing a smart phone or PDA with a workstation has the added benefit of creating a backup on the workstation's disk. Security managers would do well to circulate an occasional reminder for users to synchronize or back up their PDAs and smart phones to prevent data loss should they lose or damage their valuable tool. USB or FireWire cable connections as well as Bluetooth links may be available on different phones. Some phone or PDA docking cradles have a prominent button that activates instant synchronization, which may be completed in only a minute or two (depending on what type of data flow is configured).

One other option for backing up smart phones is cloud-based repositories. For example, Android phones can upload their configuration to and from Gmail accounts.

57.3.4 Testing. Modern backup software can automatically verify the readability of backups as it writes them; this function should not be turned off in a misguided attempt to shorten the backup time.

As mentioned earlier, when preparing for any operation that destroys or may destroy the original data, two independent backups of critical data should be made, one of which may serve as the new original; it is unlikely that exactly the same error will occur in both copies of the backup. Such risky activities include partitioning disk drives, physical repair of systems, moving disk drives from one slot or system to another, and installation of new versions of the operating system.

57.3.5 Known-Good Boot Media. Production systems must take special care to create backup boot media that are known to be uncontaminated by malware or damage from faulty software. Therefore, immediately after the initial installation of a new operating system and its required patches, operators should create a *known-good boot medium* before any other changes are made to the system.

Before applying new patches to the operating system or installing a new program, the operators reboot the system from the previous known-good boot medium and then create the next generation once the patches or program have been installed.

57.4 DATA LIFE CYCLE MANAGEMENT. Data management has become increasingly important in the United States with electronic discovery and with the passage of laws, such as the Gramm Leach Bliley Act (GLBA), Health Insurance Portability and Accountability Act (HIPAA), and the Sarbanes-Oxley Act (SOX), that regulate how organizations must deal with particular types of data. (For more on these regulations, see Chapter 64 in this *Handbook.*) Data life cycle management (DLM) is a comprehensive approach to managing an organization's data, involving procedures and practices as well as applications. DLM is a policy-based approach to managing the flow of an information system's data throughout its life cycle, from creation and initial storage, to the time when it becomes obsolete and is destroyed. DLM products automate the processes involved, typically organizing data into separate tiers according to specified policies, and automating data migration from one tier to another, based on those policies. As a rule, newer data, and data that must be accessed more frequently, are stored on faster but more expensive storage media, while less critical data are stored on cheaper but slower media.

Having created a backup set, what should be done with it? And how long should the backups be kept? This section looks at issues of archive management and retention from a policy perspective. Section 57.5 looks at the issues of physical storage of backup

media, and Section 57.6 reviews policies and techniques for disposing of discarded backup media.

57.4.1 Retention Policies. One of the obvious reasons to make backup copies is to recover from damage to files. With the 2006 changes to the Federal Rules of Civil Procedure (FRCP), the corporate legal staff should provide advice regarding the retention period of certain data in order to comply with legal and regulatory requirements and to avoid potential litigation. In all cases, the combination of business and legal requirements necessitates consultation outside the information technology (IT) department; decisions on data retention policies must involve more than technical resources. Backup data may be used to retrieve files needed for regulatory and business requirements; the next section looks at archive systems as a more effective method of storing and retrieving data needed for business and legal purposes.

The probability that a backup will be useful declines with time. The backup from yesterday is more likely to be needed, and more valuable, than the same kind of backup from last week or last month. Yet each backup contains copies of files that were changed in the period covered by that backup but that may have been deleted since the backup was made. Data center policies on retention vary because of perceived needs and experience as well as in response to business and legal demands. The next sample policy illustrates some of the possibilities in creating retention policies:

- Keep daily backups for one month.
- Keep end-of-week backups for three months.

If your company does not have a data archive system, or if the backup system does not include archival features, you may need to keep the backup media for longer periods:

- Keep end-of-month backups for five years.
- Keep end-of-year backups for 10 years.

With such a policy in place, after one year there will be 55 backups in the system. After five years, there will be 108, and after 10 years, 113 backups will be circulating. Proper labeling, adequate storage space, and stringent controls are necessary to ensure the availability of any required backups. The 2006 changes to the FRCP may require the corporate legal staff to advise retention of certain data for even longer periods, as support for claims of patent rights, or if litigation is envisaged. In all cases, the combination of business and legal requirements necessitates consultation outside the IT department; decisions on data retention policies must involve more than technical resources and may include engineering, accounting, human resources, and so on.

Readers must note that these are merely examples, not fixed recommendations on the retention periods for any organization—those must be tailored to actual requirements.

57.4.2 Rotation. Reusing backup volumes makes economic and functional sense. In general, when planning a backup strategy, different types of backups may be kept for different lengths of time. To ensure even wear on media, volumes should be labeled with the date on which they are returned to a storage area of available media and should be used in order of first in, first out. Backup volumes destined for longer

retention should use newer media. An expiry date should be stamped on all tapes when they are acquired so that operations staff will know when to discard outdated media.

57.4.3 Media Longevity and Technology Changes. For short-term storage, there is no problem ensuring that stored information will be usable. Even if a software upgrade changes file formats, the previous versions are usually readable. In one year, technological changes such as new storage formats will not make older formats unreadable.

Over the medium term—up to five years—difficulties of compatibility do increase, although not catastrophically. There are certainly plenty of 5-year-old systems still in use, and it is unlikely that this level of technological inertia will be seriously reduced in the future.

Over the longer term, however, there are serious problems to overcome in maintaining the availability of electronic records. During the last 10 to 20 years, certain forms of storage have become essentially unusable. As an example, the AES company was a powerful force in the dedicated word processor market in the 1970s; 8-inch disks held dozens or hundreds of pages of text and could be read in almost any office in North America. By the late 1980s, however, AES had succumbed to word processing packages running on general-purpose computers; by 1990, the last company supporting AES equipment closed its doors. Today it would be extremely difficult to find the equipment for reading AES diskettes.

The problems of obsolescence include data degradation, software incompatibilities, and hardware incompatibilities.

57.4.3.1 Media Degradation. Magnetic media degrade over time. Over a period of a few years, thermal disruption of magnetic domains gradually blurs the boundaries of the magnetized areas, making it harder for I/O devices to distinguish between the domains representing 1s and those representing 0s. These problems affect tapes, diskettes, and magnetic disks and cause increasing parity errors. Specialized equipment and software can compensate for these errors and recover most of the data on such old media.

Tape media suffer from an additional source of degradation: The metal oxide becomes friable and begins to flake off the Mylar backing. Such losses are unrecoverable. They occur within a few years in media stored under inadequate environmental controls and within five to 10 years for properly maintained media. Regular regeneration by copying the data before the underlying medium disintegrates prevents data loss.

Optical disks, which use laser beams to etch bubbles in the substrate, are much more stable than magnetic media. Current estimates are that CD-ROMs and CD-RW and DVD disks will remain readable, in theory, for at least a decade, and probably longer. However, they will remain readable in practice if, and only if, future optical storage systems include backward compatibility.

USB flash drives have largely replaced floppy diskettes and zip disks as a means of storing or transporting temporary data. Flash drives, similar to flash memory, have a finite number of cycles of writing and erasing before slowing down and failing. Typically, this is not an issue, as the number of cycles is in the hundreds of thousands.[35]

57.4.3.2 Software Changes. Software incompatibilities include the application software and the operating system.

The data may be readable, but will they be usable? Manufacturers provide backward compatibility, but there are limits. For example, Microsoft Word 2007 can convert files

from earlier versions of Word—but only back to version 6 for Windows. Over time, application programs evolve and drop support of the earliest data formats. Database programs, e-mail, spreadsheets—all of tomorrow's versions may have trouble interpreting today's data files correctly.

In any case, all conversions raise the possibility of data loss since new formats are not necessarily supersets of old formats. For example, in 1972, RUNOFF text files on mainframe systems included instructions to pause a daisy-wheel impact printer so the operator could change daisy wheels—but there was no requirement to document the desired daisy wheel. The operator made the choice. What would document conversion do with that instruction?

Even operating systems evolve. Programs intended for Windows 3.11 of the early 1990s would not necessarily function on Windows 7 or 8 in 2013. Many older operating systems are no longer supported and do not even run on today's hardware.

Finally, even hardware eventually becomes impossible to maintain. As mentioned, it would be extremely difficult to retrieve and interpret data from word processing equipment from even 20 years ago. No one outside museums or hobbyists can read an 800 bpi 9-track 3/4-inch magnetic tape from the very popular 1980 HP3000 Series III minicomputer. Over time, even such parameters as data encoding standards (e.g., Binary Coded Decimal [BCD], Extended Binary Coded Decimal [EBCDIC], and American Standard Code for Information Interchange [ASCII]) may change, making obsolete equipment difficult to use even if they can be located.

The most robust method developed to date for long-term storage of data is COM (Computer Output to Microfilm). Documents are printed to microfilm, appearing exactly as if they had been printed on paper and then microphotographed. Storage densities are high, storage costs are low, and, if necessary, the images can be read with a source of light and a simple lens or converted to the current machine-readable form using optical character recognition (OCR) software.

57.5 SAFEGUARDING BACKUPS.

Where and how backup volumes are stored affects their longevity, accessibility, and usability for legal purposes.

57.5.1 Environmental Protection.

Magnetic and optical media can be damaged by dust, mold, condensation, freezing, and excessive heat. All locations considered for storage of backup media should conform to the media manufacturer's environmental tolerances; typical values are 40 to 60 percent humidity and temperatures of about 50 to 75° F (about 10 to 25° C). In addition, magnetic media should not be stacked horizontally in piles; the housings of these devices are not built to withstand much pressure, so large stacks can cause damaging contact between the protective shell and the data storage surface. Electromagnetic pulses and magnetic fields are also harmful to magnetic backup media; mobile phones, both wireless and cellular, should be kept away from magnetic media. If degaussers are used to render data more difficult to read before discarding media (see Section 57.6), these devices should never be allowed into an area where magnetic disks or tapes are in use or stored.

57.5.2 Onsite Protection.

It is obviously unwise to keep backups in a place where they are subject to the same risks of destruction as the computer systems they are intended to protect. However, unless backups are made through telecommunications channels to a remote facility, they must spend at least some time in the same location as the systems on which they were made.

At a minimum, backup policies should stipulate that backups are to be removed to a secure, relatively distant location as soon as possible after completion. Temporary onsite storage areas that may be suitable for holding backups until they can be moved offsite include specialized fire-resistant media storage cabinets or safes, secure media storage rooms in the data center, a location on a different floor of a multifloor building, or an appropriate location in a different building of a campus. What is *not* acceptable is to store backup volumes in a cabinet right next to the computer that was backed up. Even worse is the unfortunate habit of leaving backup volumes in a disorganized heap on top of the computer from which the data were copied.

In a small office, backups should be kept in a fire-resistant safe, if possible, while waiting to take the media somewhere else.

57.5.3 Offsite Protection. As mentioned, it is normal to store backups away from the computers and buildings where the primary copies of the backed-up data reside.

57.5.3.1 Care during Transport. When sending backup media out for storage, operations staff should use lockable carrying cases designed for the specific media to be transported. If external firms, specializing in data storage, pick up media, they usually supply such cases as part of a contract. If media are being transported by corporate staff, it is essential to explain the dangers of leaving such materials in a car: In the summer cars can get so hot that they melt the media, whereas in winter they can get so cold that the media instantly attract harmful water condensation when they are brought inside. In any case, leaving valuable data in an automobile exposes them to theft.

57.5.3.2 Homes. The obvious, but dangerous, choice for people in small offices is to send backup media to the homes of trusted employees. There are a number of problems with this storage choice:

- Although the employee may be trustworthy, members of that person's family may not be so. Especially where teenage and younger children are present, keeping an organization's backups in a private home poses serious security risks.
- Environmental conditions in homes may be incompatible with safe long-term storage of media. For example, depending on the cleaning practices of the household, storing backups in a cardboard box under the bed may expose the media to dust, insects, cats, dogs, rodents, and damage from vacuum cleaners. In addition, temperature and humidity controls may be inadequate for safe storage of magnetic media.
- Homeowner's insurance policies are unlikely to cover loss of an employer's property and surely will not cover consequential losses resulting from damage to crucial backup volumes.
- Legal requirements for a demonstrable chain of custody for corporate documentation on backup volumes will not be met if the media are left in a private home where unknown persons may have access to them.

57.5.3.3 Safes. There are no *fireproof* safes, only *fire-resistant* safes. Safes are available with different degrees of guaranteed resistance to specific temperatures commonly found in ordinary fires (those not involving arson and flame accelerants).

Sturdy, small safes of 1 or 2 cubic feet are available for use in small offices or homes; they can withstand the heat during the relatively short time required to burn a house or small building down. They can withstand a fall through one or two floors without breaking open. However, for use in taller buildings, only more expensive and better-built safes are appropriate to protect valuable data.

57.5.3.4 Banks. Banks have facilities for secured, environmentally controlled storage of valuables. Aside from the cost of renting and accessing such boxes, the main problem for backup storage in banks is that banks are open only part of the day; it is almost impossible to access backups in a safe deposit box after normal banking hours. In any case, it is impossible to rent such boxes with enough room for more than a few dozen small backup media. Banks are not usable as data repositories for any but the smallest organizations and should not be considered as a reasonable alternative even for them.

57.5.3.5 Data Vaults. Most enterprises will benefit from contracting with professional, full-time operations that specialize in maintaining archives of backup media. Some of the key features to look for in evaluating such facilities include:

- Storage areas made of concrete and steel construction to reduce risk of fire
- No storage of paper documents in the same building as magnetic or optical media storage
- Full air-conditioning, including humidity, temperature, and dust controls throughout the storage vaults
- Fire sensors and fire-retardant technology, preferably without the use of water
- Full-time security monitoring including motion detectors, guards, and tightly controlled access
- Uniformed, bonded personnel
- Full-time, 24/7/365 data pickup and delivery services
- Efficient communications with procedures for authenticating requests for changes in the lists of client personnel authorized to access archives
- Evidence of sound business planning and stability

References from customers similar in size and complexity to the inquiring enterprise will help a manager make a wise choice among alternative suppliers.

57.5.3.6 Backing Up Online Data. A related topic is the converse of online backups: backing up online services such as Gmail and AOL, where critical data reside on a server out of a user's control. For an e-mail service, use an e-mail client configured to access the service using Post Office Protocol (POP) or Internet Message Access Protocol (IMAP) (see Chapter 5 in this *Handbook*), or use scripting languages to extract data from the server to a hard disk.[36]

57.6 DISPOSAL. Before throwing out backup media containing unencrypted sensitive information, operations and security staff should ensure that the media are unreadable. This section looks at the problem of data scavenging and then recommends methods for preventing such unauthorized data recovery.

57.6.1 Scavenging. Discarded disk drives with fully readable information have repeatedly been found for sale by computer resellers, at auctions, at used-equipment exchanges, on eBay, and at flea markets and yard sales.[37] In a formal study of the problem from November 2000 through August 2002, MIT scientists Simson Garfinkel and Abhi Shelat bought 158 used disk drives from many types of sources and studied the data they found on the drives. Using special analytical tools, the scientists found a total of 75 GB of readable data. They wrote:

> With several months of work and relatively little financial expenditure, we were able to retrieve thousands of credit card numbers and extraordinarily personal information on many individuals. We believe that the lack of media reports about this problem is simply because, at this point, few people are looking to repurposed hard drives for confidential material. If sanitization practices are not significantly improved, it's only a matter of time before the confidential information on repurposed hard drives is exploited by individuals and organizations that would do us harm.[38]

Computer crime specialists have described unauthorized access to information left on discarded media as scavenging, browsing, and Dumpster®-diving (from the trademarked name of metal bins often used to collect trash outside office buildings).

Scavenging can take place within an enterprise; for example, there have been documented cases of criminals who arranged to read *scratch tapes,* used for temporary storage of data, before they were reused or erased. Often they found valuable data left by previous users. Operations policies should not allow scratch tapes, or other media containing confidential data, to be circulated; all scratch media, including backup media that are being returned to the available list, should be erased before they are put on the media rack.

Before deciding to toss potentially valuable documents or backup media into the trash can, managers should realize that in the United States, according to a U.S. Supreme Court ruling, discarded waste is not considered private property under the law. Anything that is thrown out is fair game for warrantless searches or inspection by anyone who can gain access to it, without violating laws against physical trespass. Readers in other jurisdictions should obtain legal advice on the applicable statutes.

Under these circumstances, the only reasonable protection against data theft is to make the trash unreadable.

57.6.2 Data and Media Destruction. When a file is erased or purged from a magnetic disk, most operating systems leave the information entirely or largely intact, removing only the pointers from the directory. *Unerase* utilities search the disk and reconstruct the chain of *extents* (areas of contiguous storage), usually with human intervention to verify that the data are still good.

57.6.2.1 Field Research on Data Remanence. Garfinkel and Shelat reported on the many failed attempts to destroy (*sanitize*) data on disk drives, including erasure (leaves data almost entirely intact), overwriting (good enough or even perfect but not always properly applied), physical destruction (evidently, renders what is left of the drive unusable), and degaussing (using strong magnetic fields to distort the magnetic domains into unreadability).

With regard to ordinary file-system formatting, the authors noted:

> …most operating system format commands only write a minimal disk file system; they do not rewrite the entire disk. To illustrate this assertion, we took a 10-Gbyte hard disk and filled every block with a known pattern. We then initialized a disk partition using the Windows 98 FDISK

command and formatted the disk with the format command. After each step, we examined the disk to determine the number of blocks that had been written… . Users might find these numbers discouraging: despite warnings from the operating system to the contrary, the format command overwrites barely more than 0.1 percent of the disk's data. Nevertheless, the command takes more than eight minutes to do its job on the 10-Gbyte disk—giving the impression that the computer is actually overwriting the data. In fact, the computer is attempting to read all of the drive's data so it can build a bad-block table. The only blocks that are actually written during the format process are those that correspond to the boot blocks, the root directory, the file allocation table, and a few test sectors scattered throughout the drive's surface.[39]

57.6.2.2 Operating System Formatting. Multiuser operating systems remove pointers from the disk directory (*file allocation table*) and return all sectors in a purged file to a disk's *free space map*, but the data in the original extents (sections of contiguous disk space) persist until overwritten, unless specific measures are taken to obliterate them. Over time, as the disk continues to be used, some of the deleted file data are likely to be overwritten by chance as the free extents are assigned to new files and overwritten during file initialization and further data storage to these previously used sectors.

Formatting a hard disk using Windows' *Quick Format* option deletes no data except pointers to the file labels.[40] Full formatting does overwrite all the data on the disk, but not so thoroughly that the information cannot be read.

57.6.2.3 Clearing Data. *Erase* functions are widely available through free or purchased utilities; they typically not only remove entries in the file allocation table but actually overwrite the data up to the end-of-file using pseudorandom sequences in multiple passes. Some of these utilities explicitly overwrite slack space (the sectors between the end of file and the end of the last extent).

Even full formatting and explicit overwriting files on magnetic media do not make them unreadable to sophisticated equipment and forensic software. Since information on magnetic tapes and disks resides in the difference in intensity between highly magnetized areas (1s) and less-magnetized areas (0s), writing the same 0s or 1s in all areas to be obliterated merely reduces the signal-to-noise ratio. That is, the residual magnetic fields still vary in more or less the original pattern—they are just less easily distinguished. Using highly sensitive readers, a magnetic tape or disk that has been zeroed will yield much of the original information.

One way of *clearing* data[41] on magnetic media is to overwrite using several passes of random patterns. The random patterns make it far more difficult, *but not impossible*, to extract useful information from the discarded disks and tapes. For this reason, the U.S. Government no longer allows overwriting of classified media to obliterate data remanence when moving media to a lower level of protection. Small removable thumb drives of 2 GB and Advanced Technology Attachment (ATA) disk drives of 15 GB or more can use a *secure erase* to purge data.[42] ATA drives have secure erase in their firmware. Secure erase for thumb drives can be downloaded from the University of California, San Diego (UCSD) CMRR site.[43]

Degaussers can be used to remove data remanence physically. A degausser is a device that generates a magnetic field used to sanitize magnetic media. Degaussers are rated based on the type (i.e., low energy or high energy) of magnetic media they can purge. Degaussers operate using either a strong permanent magnet or an electromagnetic coil. These devices range from simple handheld units suitable for low-volume usage to high-energy units capable of erasing magnetic media, either oxide- or metal-based, up to 1,700 Oersteds (Oe). These devices remove all traces of previously recorded

data, allowing the media to be used again. The U.S. Government has approved certain degaussers for secure erasure of classified data. Degaussing of any hard drive assembly usually destroys the drive, as degaussing also erases firmware that manages the device. There are other downsides to using degaussers.

57.6.2.4 Shredding. Another solution is physical destruction[44] of magnetic or optical backup media before they are discarded. For end user departments, operations and security staff can provide identifiable secure collection receptacles (typically black) throughout the enterprise. Discarded media can be erased or destroyed by appropriate staff on a regular schedule. Hard disks, tapes, optical disks, and floppy disks can be cut into pieces, melted with oxyacetylene torches, crushed in large compactors, and incinerated, although proper incineration requires specialized equipment to prevent atmospheric release of toxic by-products. Some commercial companies specialize in secure destruction of sensitive records and can provide bonded pickup services, or mobile destruction units, that move from enterprise to enterprise on a regular schedule and handle paper as well as magnetic and optical media.

Ensconce Data Technology, a firm specializing in secure data destruction, has argued that software overwriting alone is not trustworthy because the choice of algorithm may be inadequate and because certain portions of the drive may not be overwritten at all. Degaussing is unreliable and even dangerous; sometimes drives are damaged so that they cannot be checked by users to evaluate the completeness of data wiping—even though expert data recovery specialists may be able to extract usable data from the degaussed disks. The strong magnetic fields can also unintentionally damage other equipment. Outsourcing degaussing introduces problems of having to store drives until pickup, losing control over data, and having to trust third parties to provide trustworthy records of the data destruction.

Physical shredders capable of securely destroying magnetic hard drive disks are expensive and usually offered only by outside companies, leading to similar problems of temporary storage, relinquishing control, and dubious audit trails. Ensconce Data Technology's *Digital Shredder* is a small, portable hardware device that provides a wide range of interfaces called *personality modules* that allow a variety of disk drives to be wiped securely. The design objectives, quoting the company, were to provide:

1. Destruction of data beyond forensic recovery
2. Retention of care, custody, and control
3. Certification and defendable audit trail
4. Ease of deployment
5. Ability to recycle the drive for reuse

The unit can wipe up to three disks at once. It includes its own touch screen; offers user authentication with passwords to ensure that it is not misused by unauthorized personnel; provides positive indications through colored light-emitting diodes (LEDs) to show the current status of each bay; can format drives for a range of file systems; and can be used to reimage a drive by making bitwise copies from a master drive in one bay to a reformatted drive in another.[45]

57.7 COSTS. All data-center managers should be able to answer questions about the costs of the backups being made on their systems. Exhibit 57.2 presents the factors

Disk Costs, Variable	
a. Disks/Backup	_____
b. Purchase Cost/Disk	_____
c. Disk Cost/Backup (a*b)	_____
Time Costs, Variable	
d. Hours/Backup Disk	_____
e. Total Disks/Backup	_____
f. Total Hours/Backup (d*e)	_____
g. Operator Cost/Hour (Salary + Benefits)	_____
h. Time Cost/Backup (f*g)	_____
Total Variable Costs/Backup (f+h)	_____
Fixed Costs	
i. Monthly costs of storage space, racks, insurance, transportation, hardware rental and maintenance, software rental and maintenance, cost of time to combine daily backups into weekly, weekly into monthly, etc.	_____
j. Number of Backups/Month	_____
k. Fixed Cost/Backup (i/j)	_____
Total Variable and Fixed Costs/Backup (f+h+k)	_____
Annualized Costs	
l. Backups/Year	_____
m. Total Variable and Fixed Costs/Backup (f+h+k)	_____
n. Total Cost/Year (l*m)	_____

EXHIBIT 57.2 Calculating Costs of Backup

that should be included in a simple spreadsheet when calculating costs of backups, with disks as an example of backup media.

57.8 OPTIMIZING FREQUENCY OF BACKUPS. Suppose a manager asks the security and operations staff these questions:

- "If backups are so important that you do a daily full backup, why don't you do a full backup twice a day?"
- "If taking a daily full backup is good enough for you, why don't you save money by doing a full backup only every other day?"

To answer such questions, managers must be able to adjust the frequency of backups to the perceived risk. One of the ways of approaching a rational allocation of resources when faced with random threats is to calculate the *expected value* of a strategy. The expected value is the average gain (if it is a positive quantity) or loss (if it is negative) that participants will incur in a process that involves random events. When this technique applies to losses over an entire year, it is called the *annualized loss expectancy*. Insurance companies use this approach to balance the costs of premiums against the disbursements to customers.

For backups, the principle is summarized by this equation:

$$E(x) = P(u)^{*}C(u) - P(n)^{*}C(n)$$

where:

x = Some particular strategy, such as doing a daily full backup.

$E(x)$ = Expected value or cost of the strategy.

$P(u)$ = Probability of having to use the backup within a single day; for example, 1 chance in 1,000, or 0.001.

$C(u)$ = Money saved by not having to redo all the work that would otherwise be lost if there were no backup; for example, the cost of paying for reconstruction of the previous day's data (e.g., $9,000) + avoidance of lost business, wasted salary, and other expenses during 3 hours of downtime during reconstruction (e.g., $30,000) for a savings of $39,000 per incident when the backups are available.

$P(n)$ = Probability of not having to use the backups at all in a given day = $1 - P(u)$ = 0.999.

$C(n)$ = Cost of making and storing a daily backup that will not be used (e.g., $50).

The expected value of doing a single daily full backup using the figures used in the last example is

$$E(x) = (0.001.\$39,000) - (0.990.\$50) = \$39 - \$49.95 = -\$10.95$$

In other words, the daily full backup has an average cost of about $11 per day when the likelihood of its use is factored into the calculations. This is equivalent to a self-insurance strategy to prevent larger disasters by investing money in preventive mechanisms and measures for rapid and less expensive recovery than possible without the backups.

If one adjusts the frequency of backups, the calculated loss expectancy can be forced to zero or even to a positive number; however, no self-insurer can make a profit from loss-avoidance measures. Nonetheless, adjusting the frequency and costs of backup strategies using the suggested factors and calculation of loss expectancies can help a data center manager answer questions from management about backup strategies in a rational manner. Since no one can estimate precisely how much a disaster costs or compute precise probabilities of having to use backups for recovery, these figures can serve only as rough guidelines.

In many organizations, the volume of changes follows a seasonal pattern. For example, 80 percent of all orders taken might come in two three-month periods spaced several months apart. For example, registration for college courses occurs mostly in the autumn, with another bulge in January. Boat, swimsuit, and ski sales follow seasonal variations. Despite this obvious variability, many organizations foolishly follow the same backup schedule regardless of date. It makes sense to adjust the frequency of backups to the volatility of data, since operations can schedule more frequent backups when there are many changes and fewer when the data are relatively stable.

57.9 CONCLUDING REMARKS. Backups are an essential component of operational security, and they play a crucial role in business resumption and disaster recovery. Backup policies should be developed rationally as a function of the specific

requirements of each application, and of each computer system. Because conditions change, backup policies should be adjusted as needed to reflect changing requirements. Decisions on keeping archival copies of data should include consideration of legal requirements. Backup policies should include provisions for the full life cycle of the backup media, including acquisition, rotation, storage, and destruction.

57.10 FURTHER READING

Cosgrove, R. *Online Backup Guide for Service Providers—How to Start and Operate an Online Backup Service*. Lulu.com, 2010.

EMC Education Services. *Information Storage and Management: Storing, Managing, and Protecting Digital Information in Classic, Virtualized, and Cloud Environments*, 2nd ed. Wiley, 2012.

Kyriazis, D. *Data Intensive Storage Services for Cloud Environments*. IGI Global, 2013.

Landis, C., and D. Blacharski. *Cloud Computing Made Easy: An Easy to Understand Reference about Cloud Computing*. CreateSpace Independent Publishing Platform, 2013.

Little, D. B., and D. A. Chapa. *Implementing Backup and Recovery: The Readiness Guide for the Enterprise*. Wiley, 2003.

Mather, T., S. Kumaraswamy, and S. Latif. *Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance*. O'Reilly Media, 2009.

McGehee, S. *SQL Server Backup and Restore*. Red Gate Books, 2012.

Nelson, S. *Pro Data Backup and Recovery*. Apress, 2011.

Nielsen, L. *The Little Book of Cloud Computing, 2013 Edition: Including Coverage of Big Data Tools*. New Street Communications, 2013.

Preston, W. C. *Backup & Recovery: Inexpensive Backup Solutions for Open Systems*. O'Reilly Media, 2007.

Rhoton, J., J. De Clercq, and D. Graves. *C*loud Computing Protected: Security Assessment Handbook. Recursive, 2013.

Schilling, C. O. *Online Backups and Cloud Computing: The Stupidly Simple Way To Backup Your Digital Life*. Amazon Digital Services, 2012.

Winkler, J. R. Securing the Cloud: Cloud Computer Security Techniques and Tactics. Syngress, 2011.

57.11 NOTES

1. According to a 1998 proposal from the International Electrotechnical Commission, the preferred prefixes should be KiB for kibibytes, MiB for mebibytes, GiB for gibibytes, and TiB for tebibytes to distinguish them from the powers-of-10 notations using kilo (103), mega (106), giga (109), and tera (1012), but this suggestion has not yet been accepted by the technical community. See NIST Reference on Constants, Units and Uncertainty. http://physics.nist.gov/cuu/Units/binary.html
2. SNIA CDP SIG, Continuous Data Protection Special Interest Group (2008), www.snia.org/forums/dmf/programs/data_protect_init/cdp (URL inactive). For information on Data Protection and Capacity Optimization Committee, www.snia.org/forums/dpco
3. For more extensive discussion of RAID levels, see IBM Technote, "About RAID levels and ClearCase" (2009), http://www-01.ibm.com/support/docview.wss?uid=swg21149421

4. 3D2F.com Software Directory, "Automatic Folder File Free Software" (2008), http://3d2f.com/tags/software/automatic/mirror/folder/file
5. 3D2F NFTP download, http://3d2f.com/programs/5-339-nftp-download.shtml
6. M. Wise, "Windows XP Systems Restore," Windows & .NET Magazine (September 2002), http://technet.microsoft.com/en-us/library/bb490854.aspx
7. Microsoft TechNet, "How to Roll Back a Device Driver, for the IT Pro" (2001), http://technet.microsoft.com/en-us/library/bb457021.aspx
8. In 1980, an HP7925 120 MB drive the size of a washing machine cost $25,000—about $213,333 per GB in 1980 dollars. In 2008, a Seagate 1 TB drive the size of a paperback book cost $237—about $0.23 per GB in 2008 dollars or roughly $0.07 in 1980 dollars. The compounded annual rate of decline in the cost of disk storage over 28 years was thus about 0.58 per year—pretty close to Moore's Law. See Intel Museum, "Moore's Law" (2008); www.intel.com/museum/archives/history_docs/mooreslaw.htm
9. Bit-for-bit copies are used in forensic analysis (see Chapter 55 in this Handbook) and include the entire contents of a storage medium, including control information and unused areas, in exactly the original layout found on the source device. In contrast, the term "clone" implies that all functional data are replicated but does not require identical distribution of those data on the physical media. Thus a bit-for-bit copy will include bad sectors and deleted files, but a clone will not.
10. hVault Holographic Storage Systems, "How It Works," http://hvault.com/a-brief-description-of-how-holographic-storage-works
11. hVault Holographic Storage Systems, "Products," http://hvault.com/solutions-for-any-size-application
12. SNIA, "Virtual Tape Library Buyer's Guide" (2006), www.snia.org/sites/default/files/SNIA_DMF_V6_DPI_Guide_4WEB.pdf
13. FalconStor, "Virtual Tape Library with Deduplication for the Fastest, Most Scalable Backup," www.falconstor.com/products/virtual-tape-library
14. P. Vettiger, et al., "The 'Millipede'—More than One Thousand Tips for Future AFM Data Storage," *IBM Journal of Research and Development* 44, No. 3 (2000): 323–340; http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=5389137
15. M. Humprhies, "IBM Showcases Millipede Storage." *GEEK News*, March 16, 2005, www.geek.com/news/ibm-showcases-millipede-storage-559662
16. M. Gibbs, "Storage? You Want Storage? How about 360TB per Disk?" *Network World Gearhead* (2013-07-11). www.networkworld.com/community/toolshed/storage-you-want-storage-how-about-360tb-disk
17. Federal Rules of Civil Procedure, www.law.cornell.edu/rules/frcp
18. www.emc.com/archiving/centera.htm
19. E. Mendelson and M. Muchmore, "The Best Backup Software," *PC Magazine*, March 25, 2013, www.pcmag.com/article2/0,2817,2278661,00.asp
20. Cisco, "Cisco Visual Networking Index: Forecast and Methodology, 2012–2017," May 29, 2013, www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html
21. J. Brodkin, "Bandwidth Explosion: As Internet Use Soars, Can Bottlenecks Be Averted? Bandwidth Usage Is Soaring, Driven by the Proliferation of Internet-Connected Devices." *Ars|technica*, May 1, 2012, http://arstechnica.com/

business/2012/05/bandwidth-explosion-as-internet-use-soars-can-bottlenecks-be-averted

22. H. Marks, "New Cloud Backup Pricing Targets Recovery," *Network Computing*, July 10, 2013, www.networkcomputing.com/backup-recovery/new-cloud-backup-pricing-targets-recover/240158054
23. SNIA Cloud Storage Initiative, www.snia.org/forums/csi
24. SNIA CDMI Cloud Storage Standard, www.snia.org/cloud/newcontent
25. M. Ricknäs, "Amazon Web Services Moves Backups to Cloud with New Appliance: The Appliance Costs $125 per Month, While Storage Capacity and Data Transfer Is Extra," *Network World*, January 25, 2012, www.networkworld.com/news/2012/012512-amazon-web-services-moves-backups-255284.html
26. S. Bradner, "Cloud Services: Computus Interruptus: Hard To Get Work Done when Cloud Services Go Down," *Network World*, March 20, 2013, www.networkworld.com/columnists/2013/032013-bradner.html
27. G. Merrill, "Can You Comply with Court Orders for Data from the Cloud?" *Network World*, July 18, 2011, www.networkworld.com/newsletters/sec/2011/071811sec1.html
28. "Dave," "Dropping the Ball on Dropbox," M. E. Kabay's *InfoSec Perception* Website, April 19, 2013, http://resources.infosecskills.com/perception/dropping-the-ball-on-dropbox
29. J. Legendre, "Cloud Storage: The Risks and Rewards," Norwich University IS340 Term Paper, November 14, 2012, www.mekabay.com/infosecmgmt/cloud_storage.pdf
30. P. Bright, "A Step Back in Time with Window 8's File History: Microsoft Reinvents Time Machine for the Second Time," *Ars|technica*, July 10, 2012, http://arstechnica.com/information-technology/2012/07/a-step-back-in-time-with-windows-8s-file-history
31. M. E. Kabay, "Backups as an Anti-Plagiarism Tool," *Network World*, March 9, 2011, www.networkworld.com/newsletters/sec/2011/030711sec2.html
32. S. Jelley, "Information Anesthetic: Taking the Pain out of Backup," *Symantec | Connect*, January 22, 2013, www.symantec.com/connect/blogs/taking-pain-out-backup
33. SiteVault Site Backup Utility, www.site-vault.com
34. OC48 has a bandwidth of 2.49E3 Mbps, which is 2.61E9 bps; 1 TB, which is 8.80E12 bits, takes 8.80E12/2.61E9 = 3.37E3 seconds = 56.2 minutes.
35. G. Crump, "How Long Will Flash SSD Last?" *InformationWeek*, May 10, 2011, www.informationweek.com/storage/systems/how-long-will-flash-ssd-last/229403122
36. G. Trapani, "Geek to Live: Back up Gmail with Fetchmail," *Lifehacker* (2007), http://lifehacker.com/software/gmail/geek-to-live—back-up-gmail-with-fetchmail-235207.php
37. Parts of this section are taken and modified from M. E. Kabay, "Disk Data Remanence, Part 1: The Problem," and following parts. *Network World Security Strategies* (2007), www.networkworld.com/newsletters/sec/2007/0730sec2.html
38. S. L. Garfinkel and A. Shelat, "Remembrance of Data Passed: A Study of Disk Sanitization Practices," *IEEE Security & Privacy* (January/February 2003), www.myoops.org/twocw/harvard/distribution/lectures/8/articles8.pdf

39. Garfinkel and Shelat, “Remembrance of Data Passed.”
40. Windows, “Formatting Disks and Drives,” http://windows.microsoft.com/en-us/windows-vista/formatting-disks-and-drives
41. Clearing information is a level of media sanitization that would protect the confidentiality of information against a robust keyboard attack. Simple deletion of items would not suffice for clearing. Clearing must not allow information to be retrieved by data, disk, or file recovery utilities. It must be resistant to keystroke recovery attempts executed from standard input devices and from data scavenging tools. For example, overwriting is an acceptable method for clearing media. Most media can be cleared with one overwrite. (NIST DRAFT Guidelines for Media Sanitization, SP800-88 Rev. 1, September 6, 2012). http://csrc.nist.gov/publications/PubsDrafts.html#SP-800-88-Rev.%201
42. Purging information is a media sanitization process that protects the confidentiality of information against a laboratory attack. A laboratory attack would involve a threat with the resources and knowledge to use nonstandard systems to conduct data recovery attempts on media outside their normal operating environment. This type of attack involves using signal processing equipment and specially trained personnel. (NIST DRAFT Guidelines for Media Sanitization, SP800-88 Rev. 1, September 6, 2012). http://csrc.nist.gov/publications/PubsDrafts.html#SP-800-88-Rev.%201
43. Center for Magnetic Recording Research, University of San Diego, http://cmrr.ucsd.edu
44. Destruction of media is the ultimate form of sanitization. After media are destroyed, they cannot be reused as originally intended.
45. Ensconce Data Technology, “Digital Shredder,” www.deadondemand.com/products/digitalshredder

CHAPTER 58

BUSINESS CONTINUITY PLANNING

Michael Miora

58.1 INTRODUCTION 58·1
- 58.1.1 Enterprise Risks and Costs 58·2
- 58.1.2 Types of Disasters 58·4
- 58.1.3 Recovery Scenarios 58·6

58.2 DEFINING THE GOALS 58·8
- 58.2.1 Scope 58·9
- 58.2.2 Correlating Objectives to Corporate Missions and Functions 58·10
- 58.2.3 Validating Goals 58·11
- 58.2.4 Mapping Goals to Recovery Phases 58·13
- 58.2.5 Emergency Issues 58·13

58.3 PERFORMING A BUSINESS IMPACT ANALYSIS 58·14
- 58.3.1 Establishing the Scope of the Business Impact Analysis 58·14
- 58.3.2 Interview Process 58·15
- 58.3.3 Describing the Functions 58·17
- 58.3.4 Definition of Departments and Functions 58·18

58.4 BUSINESS IMPACT ANALYSIS MATRIX ANALYSIS 58·24
- 58.4.1 Listing the Functions Organizationally 58·25
- 58.4.2 Finding Cross-Department Functions 58·26
- 58.4.3 Using the Ranking Factor 58·27

58.5 JUSTIFYING THE COSTS 58·29
- 58.5.1 Quantitative Risk Model 58·29
- 58.5.2 Generalized Cost Consequence Model 58·31

58.6 PLAN PRESENTATION 58·34

58.7 CONCLUDING REMARKS 58·36

58.8 FURTHER READING 58·36

58.1 INTRODUCTION. We are in an age where businesses and governments are turning in increasing numbers to high-technology systems, and to the Internet, to gain and maintain their competitive advantage. Businesses of all types are relying on high-technology products to build, promote, sell, and deliver their wares and services—as are government, educational, and nonprofit enterprises. All of these are dependent on technology to maintain their income, image, and profitability. Business continuity planning (BCP) is the process of protecting organizations from the deleterious effects on their missions that can result from outages in information systems.

The goal of BCP is to protect the operations of the enterprise, not just the computing systems. Prudent planning is not restricted to computer or telecommunications systems

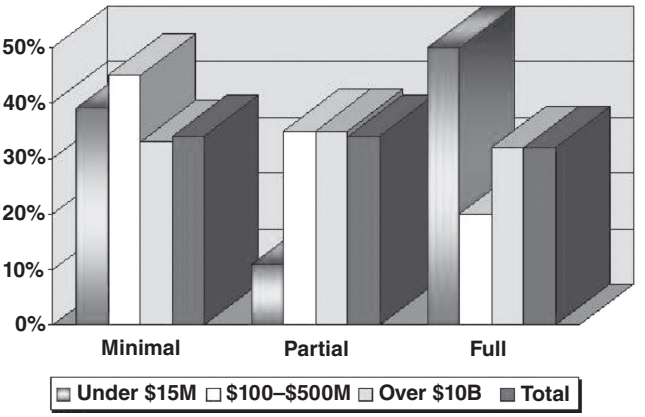


EXHIBIT 58.1 Corporate Preparedness
Source: BC Management, "2003 Benchmark Study Results."

but is enterprise wide. There lies the key difference between the traditional hot sites and service providers and the newer, self-contained, and supported backup and recovery capabilities. Without the people, procedures, and connectivity to keep the enterprise going, there is little point in restoring systems.

The enterprise has evolved to this dependency through the continual redefining of the word *fast* and the growing expectation of having access to information anywhere and at any time. In the 1960s and 1970s, fast turnaround meant a week or, with luck, a few days. In the 1980s, fast was redefined to mean *today*—the 1980s brought lower-cost computing and the beginning of high-speed communications. For the 1990s, fast meant *now*. The 1990s brought the Internet, intranets, and extranets, with distributed systems, sophisticated client/server architectures, and high-speed communications.

Today, information is predominantly collected and sent electronically, often without human intervention. Analyses have been automated and streamlined, strengthened by user-friendly tools and by expert systems. Data are requested, sent, and analyzed in minutes; reports are generated automatically; and presentations can be created, edited, and delivered in near–real time. Moreover, access to information is provided almost equally fast to internal corporate personnel, to business partners and allies, and to customers and consumers.

It should be self-evident to business owners and other stakeholders that planning is an important function and that no enterprise, large or small, should be without a plan. However, as shown in Exhibit 58.1, this is not the case. Regardless of the size of the enterprise, full readiness to handle disruptions is not universally achieved. For larger enterprises, only one-third of companies claim to be fully prepared. For smaller companies, the number is larger, but the validity of some of these responses may be questioned, given that this result is based on self-assessments reported to a survey organization.

At least one partial explanation for this lack of preparedness goes to the heart of how BCP is understood in the business community. It is not clearly understood that the goal of BCP is to protect the business, not simply its equipment or information; at least one goal of this chapter is to show that the focus of the planning should be on business issues.

58.1.1 Enterprise Risks and Costs. The overall risk is to the continued survival of the enterprise. A company weakened by a disaster and without adequate preparation may be unable to recover before failure or a hostile takeover occurs. In the

1970s, formal business continuity plans did not exist; only data backup and recovery procedures, with some disaster recovery planning for large systems, was considered. Today, companies without active and tested enterprise contingency plans are likely to fail in the event of a disaster. Such failures may have dire consequences for the board and corporate officers, whose lack of due diligence is demonstrated by the absence of planning. The need for planning may be well recognized, but protected companies are still in the minority.

There is abundant evidence that unprotected companies suffer greatly during major outages. In investment trading, time is crucial, yet many investment firms did not have a plan in place until late 1993, when a major hot-site provider first built a contingency trading facility in New York. Companies protected by disaster recovery plans can control their losses, holding them down to reasonable levels. Experts in the field believe that companies without such plans are likely to suffer significantly greater losses, with consequences so critical that almost 50 percent of firms that do not recover within 10 days never will recover.

R. A. Elbra, the author of "Contingency Planning," a technical report published in 1989 for the National Computing Center (London), demonstrated that if a critical system was inoperable without suitable backup for as few as six days, the cumulative loss could amount to 200 percent of the net daily income produced by that system. After 12 days, the cumulative loss could be 800 percent of the net daily income, as shown in Exhibit 58.2. In the years since that report was published, the reliance by large and small enterprises on their systems has grown significantly. As greater portions of corporate revenues are closely tied to systems and connectivity, these numbers take on even greater significance.

Even today, few industries have strict disaster recovery or business continuity planning regulations. The banking industry is subject to such regulations for portions of its systems only. Even Sarbanes-Oxley has minimal requirements, although there are other U.S. Federal and state laws that lay requirements on financial services organizations. The Health Insurance Portability and Accountability Act (HIPAA) contains security and privacy elements that make business continuity plans even more important than before, although full BCP is not a specific requirement.

There are many examples of serious downtime, including one instance of a construction crew repairing a nearby roadway and mistakenly damaging the power lines feeding a nearby set of buildings. Repairs required parts that delayed completion of

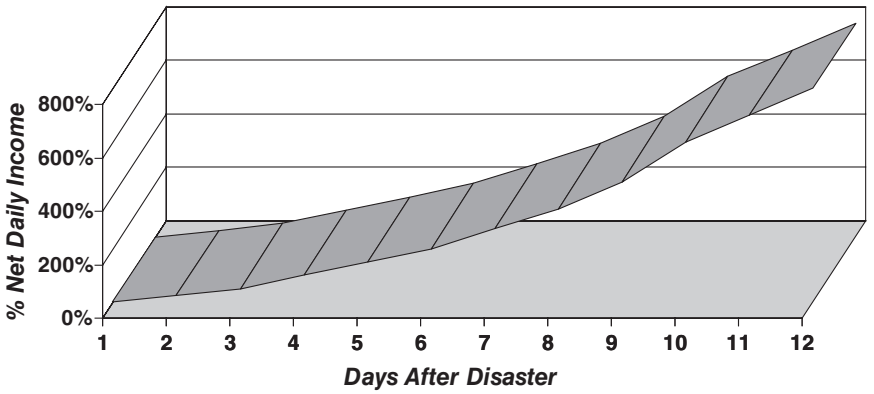


Exhibit 58.2 Fiscal Losses Caused by Disasters

repairs for almost 48 hours. One of the affected companies was an electrical supply company with revenues of approximately $50 million annually. The company felt safe because it had a plan that included reciprocal agreements to share spare computer time with other companies using similar equipment. Unfortunately, the company was unable to secure as much time as it expected because some of the other companies had changed equipment since the agreements were made, some were too busy with their own processing, and some simply did not abide by their agreements. To the chagrin of management, the replacement capability was limited to 30 percent of normal processing. Management reported that this outage cost the company over $250,000 in lost and unrecoverable sales. Although managers were unsure of their market share losses, they did know that two major customers subsequently split their orders between the affected company and another supplier. Although this is not a headline-grabbing disaster, it did significantly affect this financially healthy company.

A less healthy company might have been bankrupted by such a loss. Even surviving businesses experience significantly lower profits for the two-year period following a disaster. Reconstruction costs usually exceed insurance allowances and typically involve uninsured costs for modernization. Insurance premiums usually rise, and advertising budgets increase, to repair the corporate image. The largest problems, however, are due to productivity declines and inefficient operations caused by production restarts, debugging efforts, and the need to retrain newly hired staff to replace employees who departed in fear of corporate collapse or layoffs.

Cantor Fitzgerald, a bond company with a major office in the World Trade Center, lost 733 workers in the North Tower on September 11, 2001. The company's disaster recovery plans and mirroring sites enabled it to be back in action and taking orders just 47 hours after that loss. This was especially noteworthy because of the range of the disaster and the fact that the tragic loss of life included approximately 150 information technology workers.

Management's voice is needed to address another key issue. During natural disasters, employees at all levels of the corporate structure will make family and home their top priority. The successful recovery plan considers the need for employees with recovery responsibilities to ascertain first the condition of their own homes and families. Corporate management can make it clear to all employees that people come first. Recovery plans must provide participants with the time and the means to reach their loved ones. Once people are satisfied that families and homes are safe, then they can turn their full attention to corporate recovery.

58.1.2 Types of Disasters. Many different threats can lead to disasters. Exhibit 58.3 lists a small sampling of these potential threats. It is neither feasible nor desirable to design a strategy for each of these disasters. It is more important to consider the effects of each applicable potential threat. For example, a bomb threat may deny access to the building or to the local area for a time. A small fire may deny access to an entire building for a day and to a small portion of the building for a few months. All threats can then be grouped by their levels of impact. Exhibit 58.4 lists one hierarchical structure that maps the levels of impact to a predefined duration of outages.

Identifying disaster threats and grouping them into disaster types fulfills a planning need. Each disaster type is associated with an outage duration, which is the length of downtime expected, and with a set of predefined outage durations, such as indicated in Exhibit 58.4. Therefore, each disaster type can be mapped to specific corporate functions that will be affected.

♦ Accidents	♦ Equipment Problems
♦ Airports	♦ Fire
♦ Arson	♦ Flood
♦ Blackout	♦ Forced Evacuation
♦ Bomb Threat	♦ Hurricane
♦ Building/Facility Inaccessibility	♦ Ice Storms
♦ Chemical Spills	♦ Key Resource Loss
♦ Civil Disorders	♦ Labor Disputes
♦ Civil Disturbances	♦ Malicious Damage
♦ Computer Compromise	♦ Municipal Services Stoppages
♦ Computer Virus	♦ Postal Service Disruption
♦ Construction Damage	♦ Power Loss or Fluctuation
♦ Construction Obstruction	♦ Riot
♦ Data Diddling	♦ Security Breach
♦ Data Loss	♦ Strike
♦ Disgruntled Employee	♦ Terrorism
♦ Drug Abuse	♦ Tornado
♦ Earthquake	♦ Toxic Spills
♦ Elevator Failure	♦ Vandalism
♦ Embezzlement	♦ Volcano
♦ Environmental Control Failure	♦ Water

EXHIBIT 58.3 Disaster Threats

Area Affected	*Duration of Outage*
Partial Building	Less than 1 Day
	1–3 Days
	4–7 Days
	Week to Month
	More than Month
Full Building	Less than 1 Day
	1–3 Days
	4–7 Days
	Week to Month
	More than Month
Multi-Building (Campus Facility)	Less than 1 Day
	1–3 Days
	4–7 Days
	Week to Month
	More than Month
Local Area (Immediate)	Less than 1 Day
	1–3 Days
	4–7 Days
	Week to Month
	More than Month
Region	Less than 1 Day
	1–3 Days
	4–7 Days
	Week to Month
	More than Month

EXHIBIT 58.4 Levels of Impact and Durations

For example, a bomb threat may affect an entire building, but only for a fraction of a day. If this is so, then protecting against a bomb threat requires protecting only those corporate functions performed in that building whose survival time is one day or less. A toxic spill caused by a nearby train derailment may disrupt an entire facility, or campus of buildings, for several days. Protecting against a toxic spill requires protecting all functions whose survival time is less than several days. A regional disaster such as an earthquake, flood, or hurricane, however, disrupts operations for a longer period, and may prevent quick recovery anywhere within the affected area.

Business continuity plans need to be developed in a modular fashion so that modules can be activated depending on the disaster type. Assessing the expected duration of a disaster provides a means of determining which functions should be scheduled for recovery initiation; assessing the level of impact provides a mechanism for determining what type of recovery should be performed on each function. These are the two major factors that need to be identified promptly after a disaster: estimated duration and expected level of impact.

On September 11, 2001, there were many companies whose complete operations were destroyed. In one case, all of the staff perished. Yet, the companies that had made business continuity and disaster recovery in advance did recover, some with no discernable long-lasting impact to the client base or long-term profitability. The companies that survived and flourished had preplanned, prestaged, up-to-date hot sites, either through an external provider or internally.

The key point is that planning for specific threats is not necessary. Planning for types of outages is more efficient and more effective. It may not matter whether the building is damaged by fire, water, or terrorist attack, but the extent of damage and the duration of the outage are important, for those factors determine the types of recoveries to make. Traditionally, enterprise risks were calculated on a per-threat basis. Each threat was analyzed to determine the percentage probability of that threat occurring within a period. Those probabilities were then summed over all possible threats. The resulting figure was a representation of probability of some disaster having a significant impact on business operations over a period such as a year. The disadvantage of this method is that it misleads management into believing that the probabilities are so small as to be insignificant. A different and more streamlined analysis is more useful for the disaster recovery process. It is called the *Generalized Cost Consequence model* (GCC) and is described in Section 58.5.2.

58.1.3 Recovery Scenarios. Recovery scenarios are the planned steps to be followed when disaster strikes. These scenarios are designed based on the various threat levels of impact and durations. All recovery scenarios are built in three phases, or sequences, of activities as shown in Exhibit 58.5.

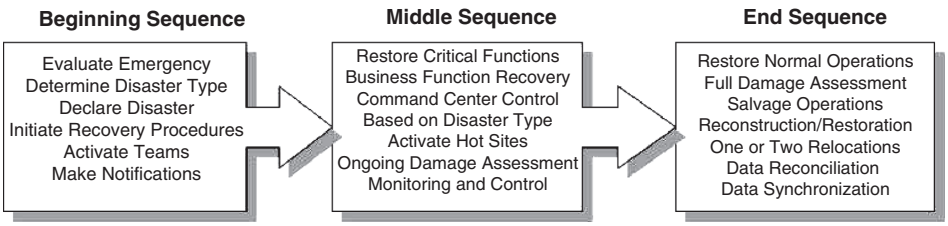


EXHIBIT 58.5 Phases of Recovery

The beginning sequence evaluates an emergency to determine the effect of the event on the enterprise. This can be an unhurried, deliberate activity when there is sufficient notice of an approaching problem. An example of this is a hurricane, where weather reports contain warnings days in advance and additional information can be evaluated as it becomes available. In contrast, the pace is highly fast and frenetic during an evacuation due to fire, flood, or earthquake. Achieving a safe position is always the first consideration; following that are the other steps: determining the disaster type and declaring a disaster; activating the disaster recovery teams; making legal, regulatory, and other notifications; and establishing command and control centers. This beginning sequence often involves top corporate management.

The middle sequence includes most of the major recovery activities. Whereas the beginning sequence determines the type of disaster in progress, this sequence provides procedures to the various functional groups based on the type of disaster in progress. One key element of the middle sequence is to perform ongoing damage assessments and maintain contact with salvage operations. Knowledgeable information technology personnel should be available to review salvage possibilities and rebuilding plans for information equipment.

The end sequence restores normal operating conditions. This phase of recovery is the least predictable of all, and is linked to disaster type and to specific threats. Therefore, this phase is planned in more general terms than the two preceding sequences. In the end sequence, detailed damage assessment is performed. Salvage operations for vital records, information equipment, data, and general facility equipment take place. Reconstruction or relocation activities are specified and executed. Finally, normal operations are restored. Often the end sequence includes two facility relocations. The first relocation moves operations from the emergency operations centers, including the command center, hot sites, work area recovery centers, and other offsite workplaces, to a longer-term interim operating area. This longer-term operating area may be a cold site or other temporary work area. The second relocation is from the interim facility back to the reconstructed, or to a newly constructed, permanent home.

Completion of the recovery plan occurs when normal operations are restored. Ending the formalized recovery procedures before then may introduce errors and may compromise integrity in the final move to the permanent facility. In this final move, large processing equipment is torn down and moved, just as it was during the initial disaster declaration, only in a more deliberate manner. Temporary operations may require use of the reserve systems as backup in case the move encounters problems. Often much of the equipment at the interim site is required at the permanent site. Finally, data reconciliation, the merging of data from different operating environments, and data synchronization require planning and careful implementation to prevent the destruction and replacement of good data with erroneous information.

The traditional disaster recovery strategies include hot sites for data centers and emergency services for employees. These sites have been effective and continue to be necessary, but not sufficient. New strategies are needed to consider the new ways of doing business: mobile computing, real-time sales and service support, client-server architectures, extensive telecommunications, real-time process controls, and online customer service support. Preplanned recovery capabilities can resurrect functions that depend on huge databases and large systems in hours rather than days. Connectivity can be restored in minutes, and the client base need never know that a disaster has occurred. Business can continue, no matter what has happened, so long as proper recovery precautions have been taken.

58.2 DEFINING THE GOALS. Defining the specific goals of a business continuity plan is a multistep process that correlates what is important to the enterprise with what the enterprise does, tempered by what is possible. A general goal for a plan may be stated in this way: Protect critical business functions so they will continue meeting minimum corporate objectives cost effectively during times of crises and disasters. This general statement describes (BCP) goals for many companies. Before building a plan, however, specific goals must be conjectured by the planner, evaluated by the planning team or steering committee, and validated by corporate management.

This section describes the specific elements and steps of defining the goal of the recovery and continuity plan, and identifying the specific objectives required to meet that goal. Succeeding sections describe how to determine the criticality and time sensitivity of various corporate business functions, how to gain and keep commitments from management, and how to define what protections are needed for each business function.

A business continuity plan may have many specific goals. For example, a central order entry and production scheduling system may reside at corporate headquarters or at the corporate sales office, while manufacturing may be dispersed geographically and by product across many facilities. One possible objective of the plan for such an enterprise is to maintain production levels at some fraction of normal levels even if the corporate center is damaged or destroyed.

A California manufacturer of custom filters devised exactly such an objective. For the vast majority of the firm's products, normal operating procedures were for a five-day cycle from order to shipment. Order entry was performed at corporate headquarters. Full material requirements planning (MRP) for all North American manufacturing sites, including purchasing, control, scheduling, and shipping, was performed on the same system. Corporate management determined that the company could not survive for more than a short time unless overall plant production capacity was maintained at 50 percent of normal levels. Production of less than 50 percent would make it infeasible for the company to survive more than two weeks. Production of more than 50 percent but less than 100 percent of normal levels would severely reduce or eliminate profits but would not cause corporate failure. The clear goal of this plan was to protect the functions required to meet the objective of achieving a minimum of 50 percent of normal levels.

The firm performed an analysis of products and revenues. It considered the quantity and cost of the various filters made, the resulting manufacturing load, and the customer base. Although the overall requirement was to maintain a manufacturing load of at least 50 percent of normal, the firm also wanted to make certain that major customers were served and that large contract commitments were met. All of these requirements were incorporated into the final BCP.

In another case, a major electronics manufacturer provided free, seven-day, 24-hour customer service and technical support to all customers. Two facilities in the United States provided support to customers worldwide. This support was a major competitive advantage and was believed to be a major contributing factor to the reputation and exceptional customer loyalty enjoyed by the manufacturer. During normal operations, the call center guidelines were:

- All incoming calls must be answered by the third ring.
- Average waiting times must be less than two minutes.
- Maximum waiting times would not exceed six minutes without operator intervention. Moreover, busy signals should not exceed one caller in 500. Communications systems linked technical support personnel with technical databases.

These were strict criteria for the company's service-level objectives, and traffic engineering studies carefully monitored busy signals and queuing delays to prevent violations of call center guidelines. Trunk lines were added or removed, and staff members were rescheduled to maintain this level of support, which far exceeded industry standards. The major goal for the call center business continuity plan was to maintain service-level objectives at a predefined minimum degradation level. The allowable queuing delay was increased significantly, and the busy signal allowance was increased slightly. A special message would play apologizing for any inconvenience, but all customers still would be served. The planning project identified what capabilities were needed to meet these requirements. The goal of this plan was to provide the backup capabilities needed to meet the reduced service-level objectives using a combination of the preexisting surviving call center and some emergency capability for the damaged or destroyed call center.

Corporate goals are independent of means and systems; they are stated in business terms. Top management must make decisions about what needs to be protected and how protected it needs to be. Often, though, this is an iterative process in which management designs some preliminary decisions and instructs analysts to confirm or refine those thoughts. The disaster recovery planning process includes steps to do just that during the business impact analysis (BIA) phase.

The goals of the plan are associated with the products and services that the enterprise or operating unit provides. In cases where the facility to be protected is a support facility, such as the corporate headquarters of a manufacturing conglomerate, the services provided are in support of the overall corporate operations, perhaps including banking and other support activities. In these cases, although the services are somewhat removed from actual production, they are still imperative to continuing production. Building the plan requires goals that are more specific in nature than those in the examples above. In this methodology, these lower-level goals are referred to as the plan objectives.

The BCP objectives are closely coupled with specific business functions and are not adequately defined until the conclusion of the business impact analysis. In order to begin this definition process, however, first the recovery problem must be constrained by defining the scope of the disaster recovery plan.

58.2.1 Scope. The scope of the plan is the definition of the environment to be protected. Before performing detailed analyses and interviewing management, decisions must be made regarding who and what is to be included in the plan. Specifying the systems, equipment, procedures, locations, and support capabilities that require protection identifies the scope. The corporate environment consists of people, information, facilities, and equipment. The plan focuses on a subset of the people, some of the information, selected facilities, and specific equipment. That constitutes the scope of the BCP. Stated more simply, no matter what the size or complexity of the enterprise, planning is performed on one part of the enterprise at a time.

The scope must be established in order to define the goals of the plan. However, during the planning process, the scope will be refined and redefined as the analysis progresses. Which facility is involved? Is part of a building being protected, such as the data center or the shop floor? Is it the entire facility, including a campus of buildings that fall into the scope of the plan? To perform a BIA, users must focus on the business elements that reside or utilize the people, information, facilities, and equipment that are in scope. However, sometimes it becomes clear during the analysis that the scope must be broadened or can be restrained to protect business functions adequately.

In a world of tight budgets, compressed schedules, and phased implementations, BCP usually is performed in stages. Perhaps the computer and telecommunications systems are protected first, followed later by other office support equipment such as copiers and hard copy files. No matter which elements are chosen for implementation, the scope must be clearly documented. For computer systems, the current systems in use must be identified. Hardware and software must be described in detail. Special equipment must be identified.

In the planning process, it sometimes becomes clear that important items have been excluded from the plan. Where those newly uncovered requirements prove to be critical, they must be included, but for all others, a well-defined scope will prevent creeping requirements that increase cost and lengthen implementation schedules. Requirements identified as desirable but not absolutely necessary can be recorded in an ongoing project log and scheduled for implementation in a later phase. It is better to put into place a plan that leaves out some capabilities while meeting basic objectives than to allow the planning scope to expand, thereby risking long delays or project abandonment. A clearly defined scope enables the planning team to communicate to management precisely what is and what is not protected or included in the planning process.

Major scope issues and disagreements sometimes can be resolved in this early phase of the planning project. If not, resolution of these issues is best delayed until the conclusion of the BIA, since that process reveals hidden functional interdependencies. For example, research and development (R&D) functions usually are longer-term projects that can be delayed without major impact on the enterprise. In many plans, R&D functions are relegated to low priority. Consider, however, the case of a customer service technical representative discussing the status of a custom-manufactured product with a key customer who is unhappy with the delivered prototype. If R&D technical personnel are required to resolve such issues, and resolution has an impact on production schedules, then R&D may be elevated to a higher priority. The BIA process usually reveals and prioritizes such hidden interdependencies.

58.2.2 Correlating Objectives to Corporate Missions and Functions.

It usually falls on the information systems and technology department to build the BCP. Naturally, then, most plans focus first on computers and telecommunications. While there is nothing intrinsically wrong with starting that way, it must be just the first step of the analysis. The danger is the potential for myopic views of corporate functions; corporate computer and telecommunications systems may be required for the successful execution of critical business functions, but usually they are only part of many needed capabilities. Desktop paper files, telephones, personal phone directories, and copying machines may be just as important as access to the corporate database or to some tailored vertical application. The list of functions and equipment to be investigated must be complete or the resulting plan will be inadequate.

The BIA provides a formal methodology for ranking business functions by criticality and time sensitivity. The process includes interviews with key personnel from each business function included in the scope of the planning project. There are two dangers inherent to an information technology–based approach to disaster recovery planning.

1. Business functions that do not use information systems may be inadvertently overlooked.
2. Noncomputer support structures and systems may not be recognized and, therefore, not brought into the protection definition process.

When the information systems department leads the disaster recovery project, it will naturally base the plan on a systems-based knowledge of the enterprise. The accounting department, a traditional heavy user of computer systems, will naturally be included among the functions that potentially need protection. Similarly, all other functions that use computers in their normal activities will likely be included in any list of functions requiring analysis. All organizational elements that use computer systems will be recognized and included precisely because they use computers and therefore are known to the systems department. However, there are many important functions that may not regularly use computers or may use them in ways that are only minimally visible to the systems people.

Mailroom operations are frequently omitted because often they are not computer users. Planning personnel are clearly well aware that mail operations exist; they just might forget to include them in the planning process even though mail operations are important to every organization. Planning mailroom recovery and continuation is simple before a disaster but very difficult after the fact.

Facilities management, maintenance, and plant engineering organizations are seldom intensive computer users in their daily activities. Nevertheless, in normal circumstances, operations in a large facility would deteriorate quickly without these functions. During a disaster, these are the people who can rapidly determine the nature and breadth of facility damage. They are the people who need to be instantly available to evaluate damage, hasten repairs, and estimate duration of outages. They are also the individuals who are often left off the disaster recovery team lists and phone rosters. Similarly, security forces are needed when building damage leaves valuable equipment and other assets vulnerable to theft or vandalism. They may be needed to protect employees who are working during civil disturbances or other such incidents. Corporate communications and public relations departments are also frequent computer systems users, but they are not very visible to the systems department as unique functions, although their contributions to a successful recovery are often essential.

The safest way to compile a list of business functions is to work from three documents: an organization chart, a corporate phone directory, and a list of corporate operations budget line items (see Exhibit 58.6). Using these three lists, the planner can develop a comprehensive list of business functions and can identify the manager or supervisor of each function. That comprehensive list should include every function performed, regardless of the perceived importance of that function. Criticality and time sensitivity will be determined during the BIA; leaving off any functions can lead to inaccurate BIA results.

Working from that comprehensive list, a second list should be compiled that includes or excludes each business function based on the defined scope of the BIA. When doubt exists, the function should be included. The BIA will be the final filter for out-of-scope functions. Therefore, for each function in the list, the planner must determine which of the included elements are relevant to the function under evaluation. Although this may seem to be a complicated process, it is usually straightforward. The golden rule is: When in doubt, include the function. It can always be excluded later if analysis shows that the function is really out of scope.

58.2.3 Validating Goals. A full disaster recovery plan must consider the contribution of each element of the organization to the overall corporate goals. Clearly, no single area is dispensable, for if it were, it would have been discarded already. Properly defined goals are independent of specific functions as they are stated in global terms. Validating the goals requires examining each function to determine how it contributes

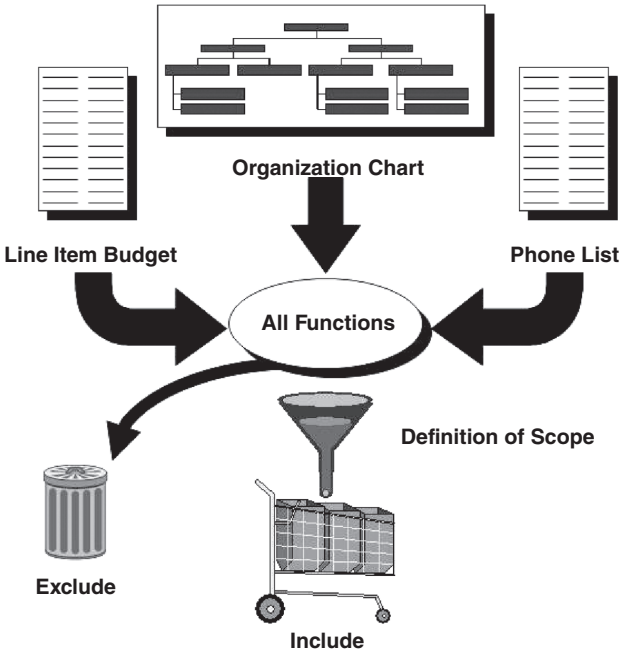


EXHIBIT 58.6 Building and Filtering the List of Functions

to the corporate objectives. It is necessary to distinguish between the importance of a function toward meeting specific disaster recovery goals and the overall importance of a function to an enterprise.

Even functions that are not relevant to any specific disaster recovery goals may be crucial to the overall operation of the enterprise. The BCP focuses on protecting selected corporate functions for a specified period of time, but it is a mistake to assume that a function whose importance to recovery goals is low is not important to the enterprise in the longer run. For example, during a crisis, the corporate tax accounting function often can cease operations with little impact to the enterprise. If this is true, then the continuation requirements for that department are minimal, and the recovery timeline can be extended. Clearly this does not mean that the corporate tax accounting function is unimportant. Other examples of longer-term requirements include fulfilling regulatory and legal requirements.

The goals of the BCP are associated with the products and services of the enterprise and are expressed in business terms. Validating those goals requires presenting management with options for levels of protection and their associated costs. Gross estimates are possible at the outset; refined presentations require a complete BIA. The goals should be clearly defined and presented to management. Although these goals are not addressed specifically by each procedure in the ultimate BCP, they drive the entire recovery strategy development. A clear and concise description of the goals must be presented to management and approved by them. After this approval, the disaster recovery planners can attempt to attain those goals. During the strategy development, other alternatives may appear and may be added for ultimate management consideration, but the initial goals become the operational baseline against which strategies and costs can be measured.

58.2.4 Mapping Goals to Recovery Phases. There are three phases to the recovery process:

1. Continuation activities
2. Resumption activities
3. Restoration activities

The continuation activities are those specific tasks and procedures that enable a very limited set of functions to continue operating with little or no interruption. The resumption activities provide for resuming a full, or almost full, range of business functions, even if that resumption means using backup methods and temporary operating procedures. The restoration activities are those that bring back a normal operating environment in a permanent facility.

The continuation and resumption activities occur during the middle sequence of activities and are the activities that must meet the short-term recovery goals that are within the domain of the recovery plan. The restoration activities occur during the end sequence and may include some long-term goals. The long-term goals may include decisions about rebuilding versus relocating and other major decisions that are beyond the province of BCP construction.

Each goal should be assigned to one of these three sets of activities. This assignment is based primarily on timeline considerations: what must be continued quickly versus what can wait a short while, and what can be delayed for the longer term. In the call center example, the number of rings, the allowable queuing delays, and the busy signal allowance were measurements of service levels. They constituted the goals for the recovery plan. After an interruption, the levels of service can be restored to normalcy gradually, as a function of time since the disaster. Stated differently, with each passing day, the level of service should improve. Therefore, the goal of the continuation activities may be for service within some stated parameters, and the goal of the resumption activities would be for improved levels of service. Normal levels may await the restoration activities.

Clearly, the same level of service can be attained at lower cost if the timing requirements are loosened. That is the reason it is important to assign each goal to a particular set of activities. Exhibit 58.7 illustrates a potential mapping of service levels to activity sets. In the first hours following a disaster, the continuation activities maintain a minimum required level of functionality, with some increases over time. The resumption activities dramatically increase functionality immediately and then bring functionality, over time, to an almost normal condition. The restoration function, shown as two bars to represent temporary and permanent facilities, achieves full levels of service.

58.2.5 Emergency Issues. A good plan includes provisions for ensuring the *safety* of all employees potentially affected by a disaster. This is especially true for employees working on site when a disaster occurs. Life safety issues include health and safety preparedness, shelter and care of employees caught in the enterprise facilities during a disaster, and search and rescue teams for employees in imminent danger. Some of these protections are required by various local, state, and federal regulatory agencies. Others are not mandated but are nonetheless important to the survival of the employees and the business.

Public relations is an important issue during any disaster. More than one enterprise has realized too late that neglecting the media can induce a second disaster, even

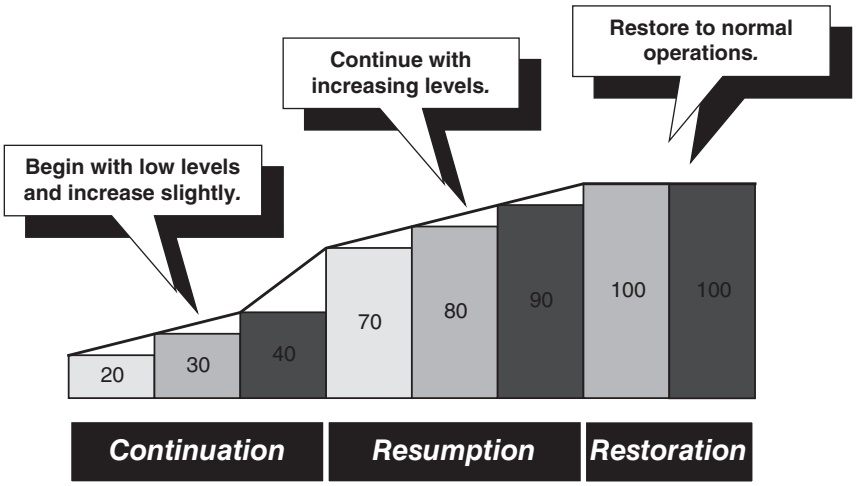


Exhibit 58.7 Service Levels Mapped to Recovery Activities

greater than the original one to which the media was responding. Even a small fire can make major news in the local community. The public relations aspect includes recovery actions. For example, one chemical company suffered a minor fire but refused to inform the media on the status of cleanup operations. The result was predictable: The media reported that a fire had potentially released toxic chemicals into the local environment and that the company had refused to comment. That report did more to damage the reputation of this chemical company than did the fire that caused the spill, which was not in fact toxic and was brought under control within hours of the fire.

58.3 PERFORMING A BUSINESS IMPACT ANALYSIS. The business impact analysis comprises the heart of the planning process. It is here that the recovery planner determines what is important for inclusion into the BCP and what is not relevant to that effort. The BIA determines how far to go in protecting the people, information, and equipment that constitute the organization and its functions so that all survive to flourish another day.

The BIA assesses how unavailability of each system would affect the enterprise. The BIA is a multistep process that may be performed over a period of weeks or months, depending on the availability of various corporate personnel. Often, the BIA may be a first effort to determine the total cost and the schedule required for a full plan. In all cases, the BIA should precede any other planning activities since it will help determine the direction and strategies for prevention, mitigation, and recovery.

58.3.1 Establishing the Scope of the Business Impact Analysis. To succeed, a project must be well defined in terms of work effort and work product. Just as the BIA helps bound the recovery problem in precise recovery terms so that the project may be successful, the BIA subproject also must be understood in terms of the work required and the product of the analysis. This requires establishing the scope of the BIA by deciding what equipment the BIA will investigate and what people to interview.

The BIA begins with an inventory process to catalog the various equipment and capabilities to be protected. As systems evolve to meet ever-increasing customer demands,

equipment and management responsibilities change along with them. For example, a mainframe shop with large local area networks (LANs) and wide area networks (WANs) may be managed by the information technology department, whereas smaller LANs may be managed by individual departments such as accounting or engineering.

Without a comprehensive and constantly updated list of systems across departments and functional areas, systems can evolve quickly with little or no overall corporate knowledge. For the recovery planner, this means the plan will be incorrect at its outset, or will grow obsolete quickly.

The inventory should describe in detail the general hardware and software used in each included system, paying particular attention to special-purpose equipment, such as imaging equipment, custom equipment, and uncommon equipment.

The inventory must include communications links and equipment. A connectivity map showing LAN and WAN equipment and connections is useful; a cabling diagram with gateways, routers, bridges, firewalls, proxy servers, and other communication equipment is also important. General telephony diagrams and summaries of lines and capacities are required to bring voice communications systems into the recovery process.

The inventory provides an opportunity to decide what level of office equipment and supplies (e.g., paper, writing implements, staplers, note pads) should be protected by the recovery plan. Although staplers seldom will be critical, certain printed business forms might well require offsite storage for quick recovery. Boxes of corporate checks, for example, should be stored securely off site.

Security access controls and special alarm systems should be included in the inventory, along with current procedures for maintaining them. There have been instances in which recovery was hampered by inaccessibility to the building, caused by active alarm and access control systems. Inclusion of such systems into the inventory helps ensure that they will help rather than hinder the recovery process.

58.3.2 Interview Process. The best source of information about the work performed by the enterprise is the enterprise's own labor force. Corporate executives understand and control corporate goals; managers understand and control operations. Front-line supervisors and workers perform the daily tasks that bring revenue to the enterprise, and these are the people who have the knowledge needed for the BIA. The recovery planner must collect that information, understand it, and translate it into terms meaningful for disaster recovery planning. Exhibit 58.8 shows the three steps of interviewing.

The first step is to compile a list of all the departments that fall within the scope of the BIA. The term *department* is used here to describe an organizational entity that may not correspond with a specific organization chart position or with the use of the term within the enterprise. The term is meant to convey a functionally complete unit performing a task or a series of related tasks. It might well be that people from different organization chart departments work more closely with one another than they do with people in their own departments.

For example, the accounting department of an organization may handle all corporate accounting functions, including accounts receivable (A/R), accounts payable (A/P), general ledger (G/L), and payroll. However, for the purpose of the planning project, the planner may choose to group accounting into two departments: accounting and payroll. In this way, payroll can be separated from its organizational component and combined with other portions of the enterprise that provide the payroll data, with which the payroll department interacts daily.

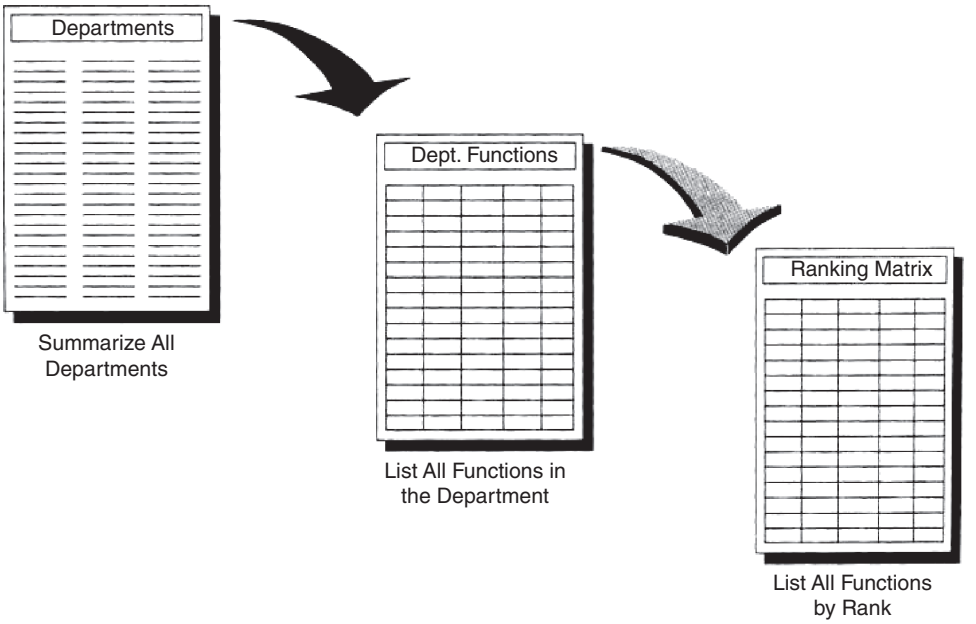


EXHIBIT 58.8 Three Steps of Interviewing

Once the list of departments is complete, the next step is to choose an individual in each department to be the primary interview candidate. Because the primary interviewee from each department will perform two functions—(1) identify all of the high-level functions performed within that department and (2) describe in detail many of those functions—that person should be an experienced and knowledgeable member of that department. It is neither necessary nor desirable to assign this interview position to a senior manager. Management interviews should be scheduled for a later phase that requires management insight. The first interviews should focus on daily task structures and purposes.

The recovery planner guides the interview. The first step of the interview is to describe the department in terms of its overall function and to list all the high-level functions performed. For example, in the accounting department, the high-level functions usually include accounts payable, accounts receivable, general ledger, and payroll. It may include corporate tax preparation, or that may be a completely separate function. It may include financial reporting for the Securities and Exchange Commission (SEC) and other filings, or a different department may perform those functions.

A corporate human resources department may include employee records, benefits management, payroll distribution, time sheet processing, and other common functions. For the planning project, the payroll distribution and time sheet processing functions may be more closely allied with the accounting department than with the human resources department.

During the interview process, the department expert will likely want to describe functions at a detailed level. For example, in the financial reporting area, the expert may want to list each and every report produced. While it may be useful for the planner to learn about the production of these reports, the BIA should categorize all related reporting functions into one function. Therefore, there is no function called *10K Reports*, but there may be an overall SEC Reporting function. It is not possible

Common Functions
♦ Accounts Payble
♦ Accounts Receivable
♦ Cash Accounting
♦ Customer Service
♦ Financial Analysis
♦ Financial Consolidation
♦ Financial Reporting
♦ General Ledger
♦ Inventory
♦ Invoicing
♦ Marketing Support
♦ Order Entry
♦ Parts Inventory
♦ Payroll
♦ POS, ATM and Related
♦ Process Control
♦ Purchase Order Generation
♦ SEC Reporting
♦ Tax
♦ Vital Records
♦ Work-In-Progress

EXHIBIT 58.9 Examples of High-Level Functions

to list all functions or potential functions here, but Exhibit 58.9 lists some examples. Each department should include several to a dozen functions. If a department contains only one function, it may be combined with another department, or it may be divided into lower-levels functions. If a department includes too many functions, then either the functions should be combined into a higher level, or that department can be divided into two departments.

The primary purpose of these interviews is to provide the information necessary to perform a matrix analysis that ranks all enterprise functions for recovery capabilities and timelines. An important secondary purpose is to raise corporate awareness of the recovery planning goals and preparations. During this interview process, the planner can explain that the overall purpose of the planning project is to protect employees and the business functions they perform. This is the opportunity to promulgate disaster preparedness as a normal and necessary part of the corporate business posture.

58.3.3 Describing the Functions. Once the functions have been listed, the interviewer must collect summary information about each function. This functional summary will be included in the BIA document, so that it will be clear to the reader precisely what the function entails. The description of no more than one or two paragraphs prevents misunderstandings and helps focus discussions during the interview. Along with the summary description, the interviewer can further describe the function by identifying the key and alternate individuals responsible for performing that

Department	Functions	Key Person	Key Alternate	Survival Days	Criticality	Ops Impact	Ranking Factor	No. Users	Category	Department Head	System Elements								

Exhibit 58.10 Blank, Sample BIA Matrix

function. These are not the managers of the functional areas; these are the people who do the work on a daily basis, enter the data, and perform the function. Each function should be associated with those individuals and with the functional area manager, and the information should be recorded in a matrix format.

This matrix format provides a single place to record all the nonnarrative information about the function. The matrix, which also will include quantitative data, will be combined with the narrative descriptions and some analysis to become the BIA document. The descriptions of each function identify the function clearly, assess the survivability factors of the function, and associate the function with the various existing corporate capabilities. The sections that follow describe the information to be collected about each function and suggest ways to encourage objectivity, even when collecting subjective quantification information. Exhibit 58.10 shows a blank, sample matrix.

58.3.4 Definition of Departments and Functions. The first column is entitled *Department,* such as Financial Reporting, Central Distribution Center, A/P, Merchandising, Planning, Distribution, and Purchasing. Although these are not necessarily equal in number or scope, they are units that need to be analyzed as groups. When unsure of organizational entities, the planner can use the organization chart as a starting point. The department title can always be modified and the functions expanded as the need arises during the development of the BIA.

The second column is the *Functions* list. The functions are the high-level groups of activities performed within the department. Each function is a group of activities related to a single purpose. In the accounting department, for example, there is an accounts payable function, which may include activities such as receiving invoices, requesting payment approval, scheduling payment, and printing checks. Although these detailed tasks are very important, they need not be described individually since they all follow a single thread and timeline. The aim of the BIA is to establish the time and processing needs; therefore, any series of tasks that occur in a closely connected chain can be addressed together. Exhibit 58.11 pictorially depicts the relationship between departments and functions.

Another example of a department is the distribution center of a warehousing operation. This may include functions commonly known as receiving, processing, shipping, quality control, and inventory control. Some of these are aggregates of lower-level tasks. Processing is a function that includes *pulls*, *moves*, and *put-aways*. If shipping, receiving, and processing are performed using the same resources according to a similar timeline, then they can be grouped as one function. Often the BIA calls this the

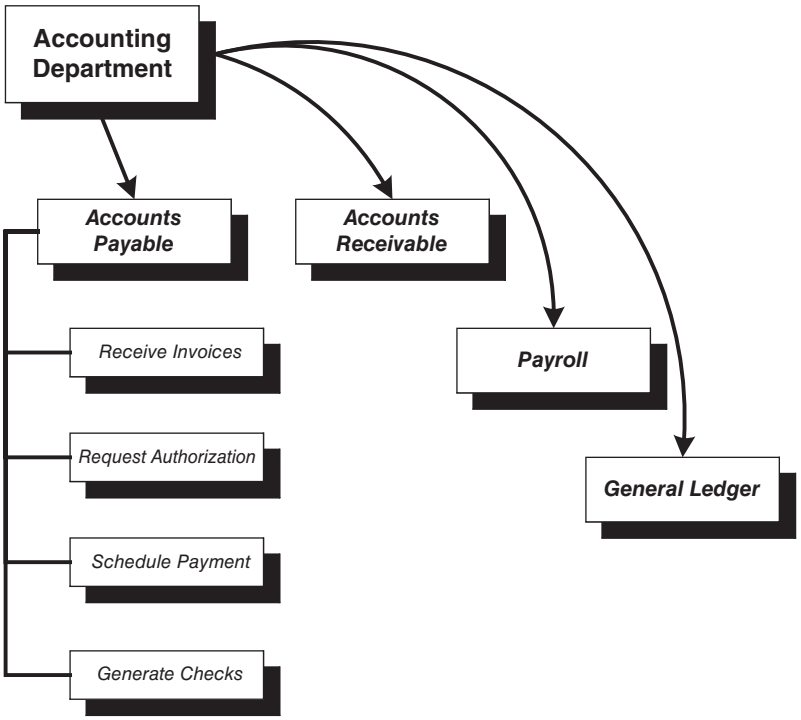


EXHIBIT 58.11 Departments and Functions

Shipping and Receiving function. If inventory control is an outcome or natural by-product of the Shipping and Receiving function, then it need not be addressed individually. If inventory control includes special functions apart from the Shipping and Receiving function, however, it can be addressed separately.

One example of inventory as a separate function is regulated inventories. For example, alcohol storage requires careful scrutiny according to regulations of the Bureau of Alcohol, Tobacco, and Firearms. Inventories must be documented according to length of storage time, source, destination, and quantities. In this case, inventory control is a special function with its own resource and timeline requirements. Another example of inventory control as a separate function is the isolated customs areas in import and export centers, where inventory must be controlled according to source, arrival time, destination, departure time, and shipper.

If functions are defined too tightly, which means at too low a level of detail, then there will be a great many functions, thereby hiding the true set of functions requiring analysis. This will result in a large set of functions that are so interrelated that they cannot reasonably be separated for individual analysis. If the functions are defined too loosely, which means at too high a level, then there will be too few functions, thereby hiding the timelines and forcing all functions to be rated as time critical. During the interview process, it is likely that the definition of functions will be refined and redefined several times. The functional area experts should guide the definitions, with the recovery specialist helping formulate the ideas.

58.3.4.1 Key Person, Key Alternate, and Department Head. The third column in Exhibit 58.10 identifies the *Key Person* for each function. The key

person is typically neither a manager nor a supervisor. Rather, the key person normally performs the task at hand. For example, in the accounting department's time-keeping function, the key person is the individual who collects and enters or validates the time-keeping information. For the treasury department, bank transfer function, this is the individual who calls or connects to the banking system to perform the transfers and manage accounts—in other words, the person who will be responsible for performing the function during a crisis when little supervision, help, or direction is available. The manager may not know the phone numbers to dial, the people to talk to, or the keystrokes required to perform the function, but the key person must know these things, for that person may be working in virtual isolation immediately following a disaster.

The key person becomes an important member of the recovery team when that function is identified as a time critical, or Category I, function. In those cases, the key person receives training, special backup and communications equipment, and extra responsibilities. The key person is the enterprise's first line of defense against disasters.

The *Key Alternate Person*, identified in the fourth column, is the backup to the key person. When possible, the key alternate should be as well equipped and well trained as the key person. In practice, however, this is seldom feasible. The key alternate should be an employee who is as familiar with the tasks to be performed as the key person.

The eleventh column identifies the *Department Head*. This is the one individual who is held accountable for performance of all the functions listed for the department. The department heads for the various departments recorded in the matrix may report at different levels on the organization chart; this is an acceptable situation so long as each department head has direct management control of all functions listed for the department.

58.3.4.2 Survival Time. The fifth column records the *Survival Days* as reported by the interviewee and discussed with the disaster recovery planner. The survival days are the length of time the enterprise can withstand the lack of a function with minimal impact. This length of time is measured in whatever units are appropriate to the enterprise, which may be minutes, hours, or months. This length of time is the maximum allowable downtime for the function, after which the enterprise begins to suffer serious repercussions. This measure does not determine the importance of the function to the enterprise; it only measures the time sensitivity of the function. Payroll, for example, is most often very time sensitive, affecting operations after only one or two days.

It may be possible to postpone tax accounting functions because there are alternative actions possible. Missing a filing deadline is avoidable by using past data and making reasonable, informal estimates of changes. Overpayments have few or no consequences. Underpayments may result in modest penalties and interest. Typically, payroll checks are generated close to the distribution date, with little spare time. However, most organizations can develop backup payroll procedures that can compress payroll processing into one or two days so that there is no delay in payroll distribution. Obtaining this information may require exploring past instances when the function had been postponed. The delay may have been due to illness, vacation, system downtime, or management direction. Inquiring about the effect of the delay may help refine the estimate of allowable down time for the BIA.

58.3.4.3 Criticality. The *Criticality* of a function, recorded in the sixth column, measures the magnitude of the immediate effect on the enterprise of function loss

beyond the survival time. If payroll processing has been determined to have a one-day survival time, then criticality is the effect on the enterprise of payroll processing delays, beginning on the second day and continuing from there. If tax accounting has been assigned an allowable downtime of 30 days, then its criticality is the affect on the enterprise after 30 days.

Criticality is a measure of time-phased impacts. The criticality of a function is loosely based on fiscal impact to the enterprise. It is not identical with fiscal loss because indirect impacts are also important. A corporate image of reliability and strength may be lost forever if functional capabilities are not restored. NASA's *Challenger* disaster had an effect far beyond any fiscal estimate. The space program suffered a setback and a long delay. Some believe that NASA's loss of funding in subsequent years was attributable to this disaster. The value of lives lost in the disaster is, of course, immeasurable.

A major telephone common carrier once lost a major switching station, thereby causing all of its customers to be without long-distance phone capability for a significant time in the middle of a business day. Although the company lost only a relatively modest amount of money due to lost calls, the damage to its reputation lives on. Its loss of business is irreparable: Many of its dedicated customers decided to double-source their long-distance providers and now split their services between this company and its major competitor. This event constituted a permanent, unrecoverable loss of market share.

The Criticality column records an estimate of the impact of loss once the survival period has been exceeded. It is an estimate because it is not the result of a detailed analysis but rather is provided by the interviewee based on heuristic analyses and experience. The interviewee is asked to rank each function on a scale of 1 to 10, where 1 is the least critical, and 10 is the most critical. Exhibit 58.12 summarizes these level criticality ratings and provides brief descriptions. The recovery planner provides this information to the interviewee and solicits a response. These criticality ratings are subjective, so the planner and the expert must work together to attempt to smooth out individual prejudices and opinions as much as possible. However, the ratings are designed to work with the operational impact measure (described in Section 58.3.4.4) to minimize this problem.

On the rating scale, a value of 10 is the highest level of criticality possible. This value should be accepted only if it is clear that virtually all corporate functions will come to a standstill. Frequently, the payroll function for hourly employees fits this criterion. Hourly, unionized employees may be specifically barred from working by union rules if the enterprise cannot provide a paycheck within some number of hours after the checks are due. In the case of a manufacturing company, this could mean that all production, receiving, shipping, and related functions cease. This is tantamount to a company shutdown. The phrase *out of business* from the summary table (Exhibit 58.12) fits this scenario. This phrase does not mean that the company ceases to exist; it means only that company operations cease until this function is restored.

Range	Description
1 to 2	"Nobody would notice." Very minor inconvenience.
3 to 4	Minor inconvenience, virtually no fiscal impact.
5 to 7	Greater inconvenience, monetary impact.
8 to 9	Major problems, significant monetary impact.
10	"Out of business."

EXHIBIT 58.12 Criticality Ratings and Descriptions

The criticality range 8 to 9 often is used to signify that loss of a function will not drive an enterprise to close its doors but will cause severe damage or loss of confidence. Missed deliveries, loss of customer service functions, lower quality, and similar effects warrant a rating in this range. Functions such as corporate communications (e.g., press releases, customer relations, and shareholder notices), banking functions, some accounting filings, and order entry are examples of functions in this range.

The range of 5 to 7 is probably the criticality rating used most often. This range indicates considerable fiscal effect and significant inconvenience. If such a function were not recovered in time, the effect would be felt but would not be a major corporate event. Perhaps the significant difference between the range of 5 to 7 and the range of 8 to 9 can be summarized in this way: The former is an internal event that would be an internal problem. The latter would be reportable in annual reports and may require media involvement or public advertisements to explain the outage. The higher range has significance beyond corporate walls. Some functions that often meet the lower criteria are accounts payable, various reporting functions, and low-priority government regulatory requirements.

The criticality range of 3 to 4 identifies loss of a function as a minor inconvenience, with little or no monetary impact. Various filings, audits, and internal support functions often can be delayed, with no impact other than the need to perform the work later when functional capability is restored. These are not unimportant functions, but they are functions that can be delayed for a fixed interval with minimal impact on corporate operations. Examples of such functions are certifications, some SEC filings, electronic data processing (EDP) audit, financial audit, internal consulting, corporate directories, and ongoing training programs. Although these are important functions, sometimes significant to overall corporate survival, often they can be delayed without major problems.

The 1- to 2-criticality level indicates that loss of the function would be barely noticeable for the short term and would cause minimal inconvenience. Although these may be important, often regulated, functions, often they are fairly time independent. Examples include former employee tracking functions, claims management, charitable contributions, and equal employment opportunity/affirmative action plans. Although these are not necessarily level 1 to 2 functions in all companies, often they are functions that can be delayed for a significant time without jeopardizing the corporate mission.

58.3.4.4 Operational Impact, Ranking Factor, and Number of Users.

The *Operational Impact*, recorded in the seventh column, is an automated result that is a function of the criticality. Individuals within the enterprise will express their biases in the criticality ratings they assign to various functions. One employee may say that a function has a criticality rating of 7 and another may say 6 or 8. The operational impact measure lowers the granularity of the estimates by transforming a scale of 10 levels to a scale of 4 levels. In this way, individual biases can be normalized and the estimates adjusted to achieve a 4-level assessment of functional criticalities and impacts.

Operational impact is derived through a many-to-one mapping that transforms criticality to operational impact, as shown in Exhibit 58.13. This transformation achieves a data-smoothing function that removes discontinuities caused by individual or organizational biases. It also reverses the order of importance, making 1 the highest impact and 4 the lowest impact. Criticality measurements are transposed; where earlier a higher number meant higher criticality to a system, now a lower number means greater impact and higher priority. This provides a mathematical convenience for combining survival time and criticality into a single measure.

Criticality	Impact	Description
10–9	1	Critical operational impact or fiscal loss
8–7–6	2	Significant operational impact or fiscal loss
5–4–3	3	Some operational impact or fiscal loss
2–3	4	No short term impacts or fiscal losses

EXHIBIT 58.13 Transformation of Criticality to Operational Impact

The eighth column is the *Ranking Factor*, which is the combination of survival time and operational impact, and hence criticality. The ranking factor is the product of the survival time in days and the operational impact. A function with short survival time and high operational impact receives the highest-priority (lowest-number) ranking factor. For example, a function with a one-day survival time and level 1 operational impact receives a ranking factor of 1. Another function with seven-day survival time and operational impact of 3 earns a ranking factor of 21.

The purpose of the ranking factor is to provide a single measure that ranks all corporate functions from highest priority to lowest priority for recovery planning purposes. The compilation of all functions listed in ascending order of ranking factor will show functions in their order of priority for disaster recovery and restoration.

The ninth column of the matrix shown in Exhibit 58.10 is the *Number of Users*. This column simply records the number of employees involved in each function. Note that in most cases, individual employees perform multiple functions and, therefore, will be counted as users in several functional areas. As a result, the sum of users in this column may exceed the total employees of the enterprise. This duplication is acceptable and accounted for in the detailed plans.

58.3.4.5 Category. The tenth column is used to record the *Category* of the function. This column summarizes the analysis achieved through the assessment of survival time, identification of criticality, and computation of ranking factors. The category is simply a way of grouping functions with similar recovery periods. Once sorted by ranking factor in ascending order, the functions are then classified into several categories based on natural groupings.

Using heuristic techniques or more formal graphical ones, the survival times of the functions can be traced or mapped as they proceed from highest priority to lowest priority. The functions will naturally fall into groups or clusters. Often one category consists of functions requiring recovery within one to three days, which can be assigned Category I, and which form the foundation of functions requiring quick recovery. Another group, which can be assigned Category II, often consists of functions requiring recovery or resumption in one to two weeks. These are the highest-priority recovery functions, once all of the Category I items have resumed. Other categories are similarly assigned.

Exhibit 58.14 shows a sample translation of functions to a graph and the subsequent assignment of categories. Each function is represented by one tick mark on the horizontal axis, the x-axis. The corresponding survival time is the height, graphed on the y-axis. In the exhibit, there are four categories: Category I functions have survival periods of from 1 to 3 days. Category II functions have survival times of 7 to 14 days. Category III functions have survival times of 30 days, and Category IV have survival times of 60 days. Although there are usually more functions than those shown on the exhibit, the process is the same.

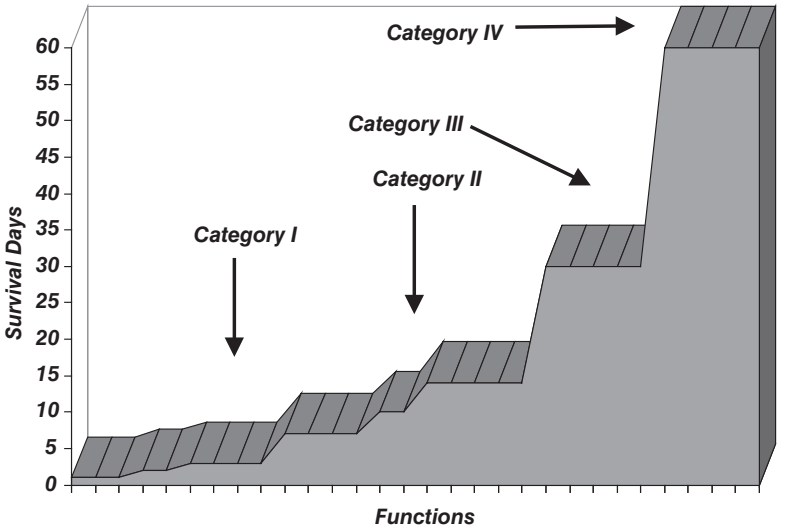


EXHIBIT 58.14 Category Assignment Curve

58.3.4.6 System Elements. The twelfth and final column is really a series of columns each representing a system element. The system elements consist of hardware, software, communications systems, and operating units. The purpose of these columns is to record which functions make use of which system elements. For each function, the disaster recovery planner marks the appropriate system element used in performing that function.

Columns may correspond to hardware devices or to a functional hardware description, such as an imaging system. Major software systems may be column headings, especially for custom-built or custom-tailored software systems, such as MRP or accounting systems. Office software systems such as word processors, spreadsheets, presentation software, databases, communication software, and many other software packages may each form column headings as well. Software packages that are always installed as a group may be described that way in the column. Several system elements may constitute one functional system. For example, the time-keeping and recording function may use electronic punch clocks, communication links between the clocks and a computer, and software to translate the electronic punch clock data into time sheets. This entire system may be labeled as the time-keeping system and be consolidated into one column.

This column also functions as a check and balance against the defined scope of the recovery plan. If columns are necessary for systems that are out of scope, then either the scope must be changed or important statements must be made regarding the limited recovery posture for those system elements not included in the plan.

58.4 BUSINESS IMPACT ANALYSIS MATRIX ANALYSIS.

The BIA matrix forms the heart of the business impact analysis. The matrix provides the basic information needed to establish recovery requirements and timelines, and to estimate costs of outages. The matrix can be manipulated to help the recovery planner perform the analyses required for translating the business objectives defined earlier in the process into the business continuity plan objectives.

58.4.1 Listing the Functions Organizationally. The major matrix manipulations that the planner will perform are a special set of sorts. Exhibit 58.15 shows a sample BIA matrix. This sample BIA matrix is representative of a corporate headquarters facility with some production and warehousing capabilities. The matrix shown is for illustrative purposes only and therefore is composed of only a selected subset of the columns found in an operational matrix. It also shows only a few of the many corporate functions that would normally comprise a full matrix.

Department	Functions	Survive Days	Criticality	Ops Impact	Ranking Factor	No. of Users	Category
Corp Accounting	Cash Management	1	10	1	1	2	
Corp Accounting	Benefits	1	5	3	3	1	
Corp Accounting	Payroll (Mgt)	2	9	1	2	2	
Corp Accounting	A/P	2	6	2	4	3	
Corp Accounting	Financial Planning	30	7	2	60	4	
Corp Accounting	Income Tax	30	7	2	60	2	
Corp Accounting	A/R	30	6	2	60	3	
Corp Accounting	G/L	30	5	3	90	2	
Corporate Services	Auto Fleet Mgt	30	6	2	60	1	
Corporate Services	Travel Administration	30	2	4	120	2	
Corporate Services	Facilities	1	10	1	1	2	
Customer Service	Order Entry	1	8	2	2	1	
Customer Service	Production Scheduling	2	9	1	2	1	
Human Resources	Payroll (Union)	1	10	1	1	2	
Human Resources	Workers Comp	10	8	2	20	1	
Human Resources	Benefits	15	3	3	45	1	
Human Resources	OSHA Compliance Programs	90	N/A	N/A	N/A	1	
Human Resources	Rideshare Program	90	1	4	360	1	
Marketing	Pricing Strategies	2	10	1	2	3	
Marketing	Strategic Product Development	14	9	1	14	5	
Marketing	Account Targeting	30	6	2	60	3	
Marketing	Advertising	90	3	3	270	2	
Materials	Production Scheduling	1	8	2	2	2	
Materials	Purchasing	1	8	2	2	3	
Materials	Bill & Materials	3	6	2	6	1	
Materials	Materials Planning	7	4	3	21	2	
Materials	Invoice I/F	14	8	2	28	2	
QA	Inspection Records	7	3	3	21	1	
QA	Project Tracking	14	4	3	42	4	
QA	Reporting	30	3	3	90	1	
R&D	CAE	1	7	2	2	4	
R&D	General Research	14	4	3	42	2	

EXHIBIT 58.15 Sample BIA Matrix

This matrix is sorted alphabetically by department. Within departments, it is sorted first by survival time of functions, starting with shortest survival. When the survival time of functions is equal, functions are sorted by criticality so that the most critical functions are listed first. This matrix sort, called the *organizational sort* because it follows corporate structures, frequently is the first one used, and is most useful for summarizing interviews and results. This sort is also the one that is most useful when meeting with department managers, for their concerns center on their own departments.

58.4.2 Finding Cross-Department Functions. A second useful view emphasizes similar functions across departmental boundaries. In this view, shown in Exhibit 58.16, functions that cross departments are grouped together. This matrix, which for simplicity is a subset of the matrix shown in Exhibit 58.15, provides facility for uncovering functions with similar titles performed by different departments. In the exhibit, there are three sets of functions with similar names but different departments.

The corporate accounting and human resources departments both perform functions named *Benefits*. However, the two departments describe the survival time and criticality of these functions very differently. Human resources considers this function a minor inconvenience with virtually no fiscal impact and an allowable downtime of more than two weeks. The accounting department tags this function with a downtime of only one day and a greater inconvenience with fiscal impact. The ranking factors of these two functions, both called *Benefits*, are widely different in value, showing at a glance that there is a discrepancy here. The planner must determine whether this discrepancy is the result of inconsistent descriptions, and therefore a problem, or the result of different functions with a similar name.

The other two sets of cross functions are the *Payroll* function, another overlap between the corporate accounting and human resources departments, and *Production Scheduling*, which is performed by both the materials department and customer service. In both cases, the survival time and criticality measures were assessed to be similar by both departments, with ranking factors that are close or the same.

Department	Functions	Survive Days	Criticality	Ops Impact	Ranking Factor	No. of Users	Category
Corporate Services	Auto Fleet Mgt	30	6	2	60	1	
Corp Accounting	Benefits	1	5	3	3	1	
Human Resources	Benefits	15	3	3	45	1	
Customer Service	Order Entry	1	8	2	2	1	
Human Resources	OSHA Compliance Programs	90	N/A	N/A	N/A	1	
Corp Accounting	Payroll (Mgt)	2	9	1	2	2	
Human Resources	Payroll (Union)	1	10	1	1	2	
Marketing	Pricing Strategies	2	10	1	2	3	
Materials	Production Scheduling	1	8	2	2	2	
Customer Service	Production Scheduling	2	9	1	2	1	
QA	Project Tracking	14	4	3	42	4	

EXHIBIT 58.16 Emphasizing Cross Departmental Functions

The matrix makes apparent the existing overlaps, but the planner must determine whether the overlaps are similar, overlapping functions, or different functions with similar names. In the first example, the function named *Benefits* is quite different. The human resources department views this function as a regular record-keeping function with occasional changes, and it includes tracking and managing employee health plan choices, usually offered annually. Short processing delays are usually inconsequential to the enterprise and its employees. For the corporate accounting department, however, the same named function includes making matching funds payments, exercising stock options, and calculating other financial disbursements based on existing benefits packages and agreements. The functions are similarly named but different in meaning.

The other two examples are similar functions with the same meanings but performed from different perspectives. *Payroll* for management and union personnel is performed differently due to union rules and regulations. Therefore, corporate accounting can calculate regular salaries and benefits for management employees and produce payroll checks. However, the human resources department must be involved in calculating pay for unionized, hourly employees to make certain that all rules and regulations are followed precisely.

Production Scheduling, performed by customer service and by materials, is a similar function with different perspectives. For customer service, production scheduling means setting requirements for customer deliveries and, thereby, determining production schedules. *Materials* finalizes the schedules to reconcile possibly conflicting production requirements for different customers by advancing or delaying schedules within corporate policy constraints.

The planner must understand the meaning of each function and recognize the differences between similarly named functions so that all recovery capabilities meet the true needs of the corporate users and customers.

58.4.3 Using the Ranking Factor. A third view, which is the most important part of the matrix analysis, is the ranking factor view. In this sort, shown in Exhibit 58.17, functions are listed in ascending order according to their ranking factor first and then in ascending sequence according to survival days. The purpose of this view is to assign *Category* designations to all functions and to establish specific timelines for the categories. These timelines become the disaster recovery plan *technical* objectives. Exhibit 58.17 consists of the same data as the previous exhibits but sorted differently.

A useful function of the ranking factor is to list functions in a reasonable priority order for recovery timeline planning. The primary purpose of the *Category* designation is to group functions by recovery requirement or timeline, as a function of short-term importance to the enterprise. Therefore, the ranking factor sort is ideal for determining functional categorizations. All Category I functions for this organization must be recovered in less than three days. For most of these functions, the criticality is high and the survival time is low. Category II functions begin when the survival time jumps to 14 days.

This type of orderly grouping is not exceptional; it occurs in most organizations. The natural organizational activity is to perform functions, generally in a regular, repeating pattern. Many functions are critical and performed almost daily. Other functions are performed weekly, monthly, quarterly, or annually. The ranking factor is a good tool for viewing functions according to their natural cycle. It works because this factor measures both cycles and relative importance. Either measure alone would be insufficient. Exhibit 58.18 graphs three basic measures: survival days, operational impact,

Department	Functions	Survive Days	Criticality	Ops Impact	Ranking Factor	No. of Users	Category
Human Resources	Payroll (Union)	1	10	1	1	2	I
Corporate Services	Facilities	1	10	1	1	2	I
Corp Accounting	Cash Management	1	10	1	1	2	I
R&D	CAE	1	7	2	2	4	I
Materials	Production Scheduling	1	8	2	2	2	I
Materials	Purchasing	1	8	2	2	3	I
Customer Service	Order Entry	1	8	2	2	1	I
Marketing	Pricing Strategies	2	10	1	2	3	I
Customer Service	Production Scheduling	2	9	1	2	1	I
Corp Accounting	Payroll (Mgt)	2	9	1	2	2	I
Corp Accounting	Benefits	1	5	3	3	1	I
Corp Accounting	A/P	2	6	2	4	3	I
Materials	Bill & Materials	3	6	2	6	1	I
Marketing	Strategic Product Development	14	9	1	14	5	II
Human Resources	Workers Comp	10	8	2	20	1	II
QA	Inspection Records	7	3	3	21	1	II
Materials	Materials Planning	7	4	3	21	2	II
Materials	Invoice I/F	14	8	2	28	2	II
R&D	General Research	14	4	3	42	2	II
QA	Project Tracking	14	4	3	42	4	II
Human Resources	Benefits	15	3	3	45	1	II
Marketing	Account Targeting	30	6	2	60	3	III
Corporate Services	Auto Fleet Mgt	30	6	2	60	1	III
Corp Accounting	Financial Planning	30	7	2	60	4	III
Corp Accounting	Income Tax	30	7	2	60	2	III
Corp Accounting	A/R	30	6	2	60	3	III
QA	Reporting	30	3	3	90	1	III
Corp Accounting	G/L	30	5	3	90	2	III
Corporate Services	Travel Administration	30	2	4	120	2	III
Marketing	Advertising	90	3	3	270	2	III
Human Resources	Rideshare Program	90	1	4	360	1	III
Human Resources	OSHA Compliance Programs	90	N/A	N/A	N/A	1	III

EXHIBIT 58.17 Ranking Factor View

and ranking factor. This graphical representation is based on the data contained in Exhibit 58.17.

Each item on the horizontal (x) axis represents one function from the matrix. The height of each of the three curves represents their numerical values. The ranking is an increasing curve because the data are presented according to the ranking factor view. The survival days is mostly increasing but does drop near the middle of the graph and then rises again. This dip corresponds to the two functions with seven-day

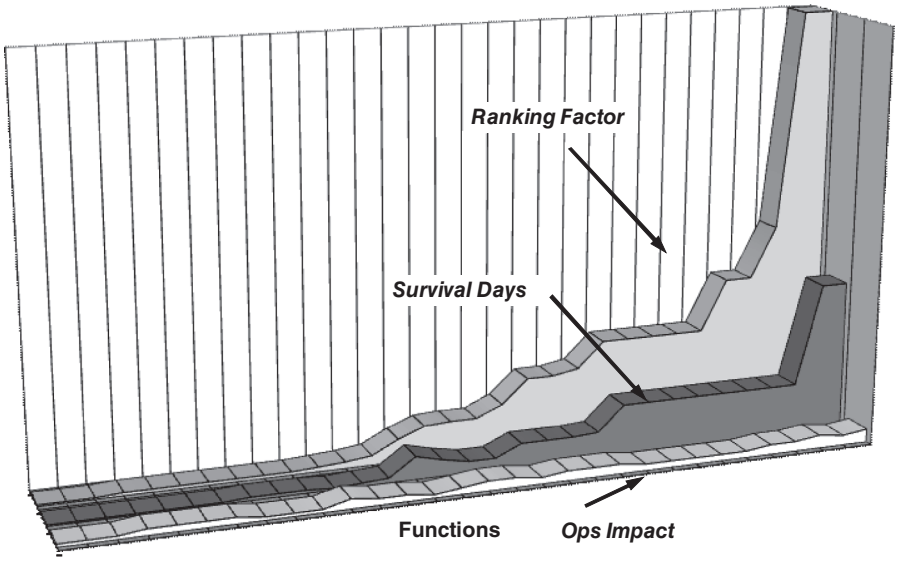


EXHIBIT 58.18 Graphing the Three Measures

survival times The rankings of the operational impact curve begin at fairly low values, corresponding to high levels of effects such as criticality 10 or 9, and end with fairly high values, corresponding to lower levels of effects, which are typically less than 5 or 6. This curve rises and falls but follows a generally increasing pattern.

Sorting this information according to operational impact would not yield sufficient information on which to base recovery timelines. Some functions with great detrimental effects only begin to have an effect after a longer time. Similarly, a simple sort of survival time may erroneously include functions as high priority whose operational impacts are too low to merit the quick recovery and expense associated with Category I functions. In both cases, the timeline would not be accurate because the sorting criteria exclude either levels of impact or time. Only the ranking factor is reliable since it reflects both time and effect simultaneously.

58.5 JUSTIFYING THE COSTS. One of the most difficult aspects of the planning process is justifying the costs of the planning process, the reserve equipment, and the contract services. The cost justification is performed best and most efficiently using the generalized cost consequence model rather than the quantitative risk model. Both models are described in the sections that follow.

58.5.1 Quantitative Risk Model. The quantitative risk model is a formal and rigorous methodology for analyzing expected losses that will be incurred over a predetermined time period. This procedure requires a significant amount of analysis and research. The model consists of three main factors: probability of loss, cost of loss, and annual loss expectancy (ALE). The probability of loss is really a sum of the probabilities of different catastrophic events that range from partial outages to severe interruptions. The cost of loss depends on the level of interruption. For example, a partial building loss affecting computer systems but leaving phone systems in operating condition has a much lower cost than a complete building destruction. Therefore, the cost of loss is dependent on the type of disaster.

A simplified risk model considers the probability of loss and the cost of the loss. The annual loss expectancy is the product of the probability and the cost. For example, let us assume there is a 5 percent probability (annual) of a major power failure. Stated differently, this means that the facility will experience a major power failure once in 20 years. Let us further assume that the power failure will cause a 72-hour outage, which will cost the enterprise $1,250,000. The ALE is calculated as 5% × $1,250,000, or $62,500. This number is compared to the baseline cost of the recovery plan and the cost of capital. For example, say the cost of the recovery plan is $2,000 per month and the cost of capital is $25,000. Then the baseline cost or comparison figure is (12 × $2,000) + $25,000, or $49,000. Since the comparison figure is lower than the ALE, recovery planning is justified.

There are some serious shortcomings in this simplified approach. First, the cost of the outage depends on the level of loss. A 72-hour power outage is significantly costlier than a 24-hour outage. Therefore, the ALE must reflect the difference in probabilities of different levels of impact. Another problem is defining the probability of occurrence for an aggregation of events.

The more acceptable risk model must consider the different levels of loss and sum the probabilities of all disasters that can cause that level of loss to define the true probability of loss for that loss level. This is accomplished for each loss level. First, a series of disaster events is defined. Each event is then refined into levels. For example, office buildings are susceptible to loss due to fire. Data on numbers of fires and amount of destruction (in predefined ranges) are available from various fire protection services. For a given facility, the total number of such buildings in the geographical vicinity is considered. Then the frequency or probability of a fire causing a range of damage is calculated. This calculation is performed for each defined range of damage. The level of loss for each range of damage is assessed, then the ALE is calculated for each range of damage by multiplying the probability of a fire causing that range of damage with cost of the loss if there is such a fire. All the ALE values are summed to calculate the total fire ALE. This calculation must be performed for all types of disasters that can affect the facility to determine the grand total ALE. For each level of impact for all disasters, the baseline costs also must be calculated. These figures are also summed to form a total baseline cost. The grand total ALE is then compared with this baseline cost figure.

This analysis is a complex process that requires great effort to generate and even greater patience to explain. However, there are two more important problems with the quantitative risk model.

1. Calculating all the outage costs is very difficult and subject to debate. Moreover, once the cost figures are finalized, they are subject to constant change due to the changing business climate and practices.
2. Calculating the probabilities is also very difficult and often requires many subjective conclusions. For example, what is the effect of modernizing the sprinkler system based on the level of damage experienced by a particular type of fire? Each countermeasure can significantly alter both the cost and the probability. Moreover, the probability of any particular event tends to be quite small—often less than 1 percent.

While the quantitative risk model is an interesting actuarial exercise, it is of marginal use in modern business continuity planning.

58.5.2 Generalized Cost Consequence Model. The generalized cost consequence (GCC) model (developed by this author between 1990 and 1995) does not consider probabilities of specific disaster events. Instead, it estimates the total cost of outages as a function of time after an event. This model is significantly simpler than the quantitative risk model: It is easier to build and simpler to explain. The GCC estimates the cost of an outage for each function and applies that cost to the total disaster cost after the maximum allowable down time has been exceeded.

For instance, assume the cost of delaying the treasury department's bank management function is $25,000 per day after the first day. Let us also assume that the cost of delaying the law department's general contract review is $5,000 per day after seven days. For the bank management function, the cost to the enterprise is calculated as $25,000 per day beginning on the second day. For the contract review function, the cost is calculated as $5,000 per day beginning on the eighth day. Therefore, the contract review function does not contribute to loss during the first seven days. This calculation is performed for each function, and then the costs are collected by category. This category cost summary is used to develop and present a graph that shows the total cost losses for each category level once they are activated and the total for all categories over time.

A sample graph of the contribution of functions aggregated by category level is shown in Exhibit 58.19. In this example, Category I functions cause slightly more than $120,000 of loss on a daily basis once the maximum allowable downtime has been exceeded. Category II functions contribute slightly under $60,000 in this example. It is likely that different functions will commence their loss contribution at various times after the disaster event. Therefore, the true Category I loss contribution may begin at a lower level and increase to its full level. That distribution of effect will occur beginning on the first day any function exceeds its allowable downtime and will continue to grow until the last day any function exceeds its downtime, at which point the effect will have achieved its full loss contribution. Since categorizations cluster functions with similar downtimes, the loss can be presented as a single, or point, value rather than as a value that varies over time.

The bars indicate that Category I has a higher value than Categories II and III. Category IV is also high. The U-shape is characteristic of this graph. No matter how many categories are used, the Category I functions tend to be quite high in their contributions. The next one or two categories are lower, but increasing in values until the lowest-priority category, which tends to be quite high. This shape results from the manner of the categorizations. The Category I functions are highly critical, with great

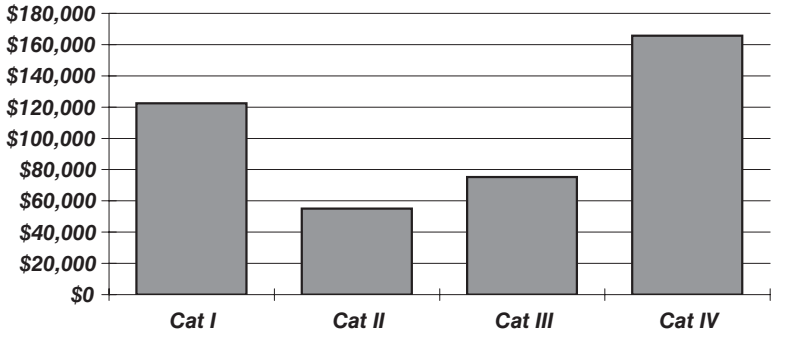


EXHIBIT 58.19 Summary of Loss Contribution

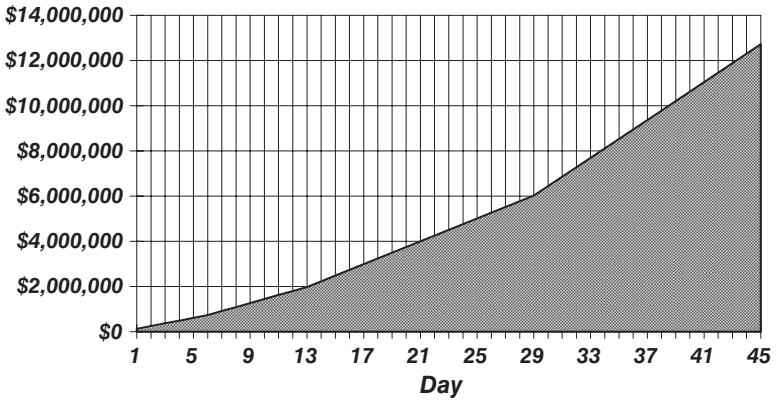


Exhibit 58.20 Cumulative Loss Summary

effects on corporate operations. Category I functions tend to be few in number, with each function contributing significantly. The middle categories tend to be larger in number than Category I, with the number increasing with each succeeding category, although still small compared to the overall set of organizational functions. The final category tends to contain the largest number of functions, each of which contributes less on a daily basis than the Category I functions, but the sheer magnitude of this set often causes it to be the largest overall contributor to loss. Of course, Category I functions begin their effects quickly, while loss of Category IV functions may be insignificant for weeks.

Exhibit 58.20 presents a sample graph showing the accumulation of losses from all categories following a disaster. This exhibit uses the same values as in Exhibit 58.19 but presents them in a different format and represents the losses accumulated after 45 days. Neither exhibit measures physical losses such as real property, capital equipment, and the like. Physical losses are not issues that affect the planning process; they affect the cost to recover and are independent of continuation measures. The horizontal (x) axis represents the number of days since the disaster event caused damage and cessation of corporate functions. The vertical (y) axis represents the total losses the enterprise will have sustained on the corresponding day represented on the x-axis.

At the time of the disaster occurrence, no losses will have accumulated. Thereafter, the accumulated loss is increased each day by the cost contribution of the category whose earliest start time has already been surpassed. Assume that Category I functions begin to contribute to corporate losses of $120,000 on the first day, that Category II functions begin to contribute losses of $60,000 on the seventh day, and that Category III functions contribute $80,000 daily beginning on the fourteenth day. In this case, the cumulative loss begins at zero and grows by $120,000 per day for the first six days. On the seventh day, the Category II functions begin to contribute $60,000 per day along with the ongoing contribution of the Category I functions. Therefore, beginning on the seventh day, the cumulative loss grows by the sum of $120,000 and $60,000, which is $180,000 daily. On the fourteenth day, the daily loss increases by another $80,000, representing the Category III contribution. This brings the total daily loss to $240,000, which is the sum of the contributions of Categories I, II, and III.

The cumulative loss summary shows at a glance the loss that an enterprise will experience over time following the disaster, if no recovery planning is performed. The active simplifications of this model are the grouping of functions by category and the

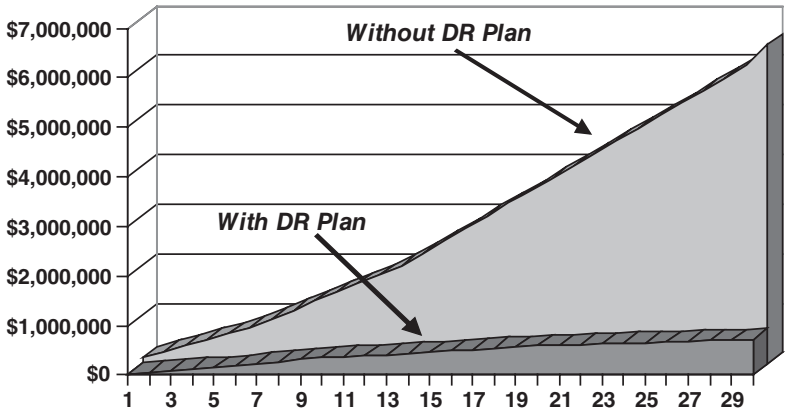


EXHIBIT 58.21 Cumulative Loss Summary With and Without Planning

subsequent representation of each category as a single value beginning at a fixed point in time. This may render the estimate slightly inaccurate at some local points, but the overall values are as accurate as the underlying estimates.

The generalized cost consequence model summarizes at a glance the effects of loss of functions. These loss figures can be affected by insurance reimbursements, legal liabilities, and overall management objectives.

A second model can be developed if the recovery planning is in place, showing the residual loss proposed. That model would be developed in a similar manner but with the assumption that certain functions are restored within established time parameters. There would be residual loss only if the restoration occurs later than the allowable downtime, which will almost certainly be the case for some functions. Exhibit 58.21 presents a graph that shows cumulative losses with and without a disaster recovery plan in place. The Without Plan curve is similar to the previous exhibit. The With Plan curve reflects the residual loss that would occur even if a plan were in place. This loss is normally dramatically lower but is seldom zero. Typically, a reduction to zero residual loss would require extraordinary and prohibitively expensive measures. Most organizations can benefit more from a substantial lowering of residual losses than from a full reduction to zero losses.

The estimation process itself is much simpler than for the quantitative risk model. For each function, the loss estimate is based on three criteria, or types of losses:

1. Tangible and direct losses
2. Tangible and indirect losses
3. Intangible losses

The tangible and direct losses are the easiest to calculate. These losses can be traced to specific revenue-producing functions. The results are direct because the loss occurs as a first-order effect, meaning that revenue stops because the function cannot be performed. The results are tangible because they can be measured easily. An example of such a function is automated production control of an assembly line. If the systems exceed their allowable downtime, then production will cease. The cost is the resultant loss of sales after inventory is depleted. Another example is loss of order entry functions. In this case, the result is similarly calculable: lost sales after in-stock

items are depleted and existing orders are produced. Tangible, direct losses include lost sales, lost manufacturing, lost deliveries, and other lost opportunities.

The tangible and indirect losses are the most common, although slightly more difficult to estimate. Support functions generally produce tangible results whose deprivation would cause a financial loss indirectly. For example, a public corporation issues quarterly earnings reports that, if late, could have significant consequences on the company's stock value. Although this is not a direct loss resulting from cessation of sales or other production, the loss can be calculated using accounting standard practices. The internal accounting personnel are in the best position to provide this estimate to the disaster recovery planner. Tangible, indirect losses include penalties, fees, fines, market share, and other issues that can be directly calculated, or at least estimated with some degree of confidence.

Intangible losses are the most difficult to calculate. The intangible effects include reduced public confidence, compromised customer satisfaction, promises not kept, damaged reputation, and other losses that are general in nature and not easily calculable. Sometimes these losses are not translated into specific financial losses and are, therefore, not represented in the cost graphs. In such cases, prominent notations should be made explaining the additional but not quantified losses.

According to a study by Contingency Planning Research in their 2001 Cost of Downtime Study, hourly losses by major companies can range from a low of $12,000 for service-oriented organizations to over $7 million for a major brokerage firm. There are no strict rules for estimating losses. The timeline can be specified in increments of days, as shown in the preceding examples. It also can be as fine as hours or minutes. Some industries, such as hospitals and other acute care facilities, must measure loss in finer increments.

The generalized cost consequence model can solve the problem of cost justification. This model shows the potential, possibly catastrophic, losses without engaging in the analysis paralysis that can stem from a detailed quantitative risk model development effort.

58.6 PLAN PRESENTATION. The development of a business continuity plan is an involved and somewhat complex process. Though there is the simplifying technique of the generalized cost consequence model, there is still significant work required to devise and implement a workable and resilient plan.

The complexity of the development process, however, must give way to simplicity in the final plan that is presented to the enterprise. Complex analyses and detailed descriptions must give way to simple and easy-to-read flows and action plans, for even after significant training, the people responsible for carrying out the plans need simple documentation to help them along.

Exhibit 58.22 provides on example of how to present the overall recovery flow. This one diagram can form the basis for a set of sections that describes each of the specific steps to be followed. Each decision point requires information and decision guidelines; the BCP document can describe each of these steps and can provide guidance for which path to follow. In this diagram, some recovery responses have *cases* that further define the steps to take. In this example, the cases are parameterized by the expected outage period or the actual outage duration.

Exhibit 58.23 illustrates the meaning of the cases in Exhibit 58.22. This illustration draws from the previously described concepts of function categories and lays those categories against their specific timeline requirements. The cases overlay time periods and make it clear which functions must be recovered and in which of the three cases

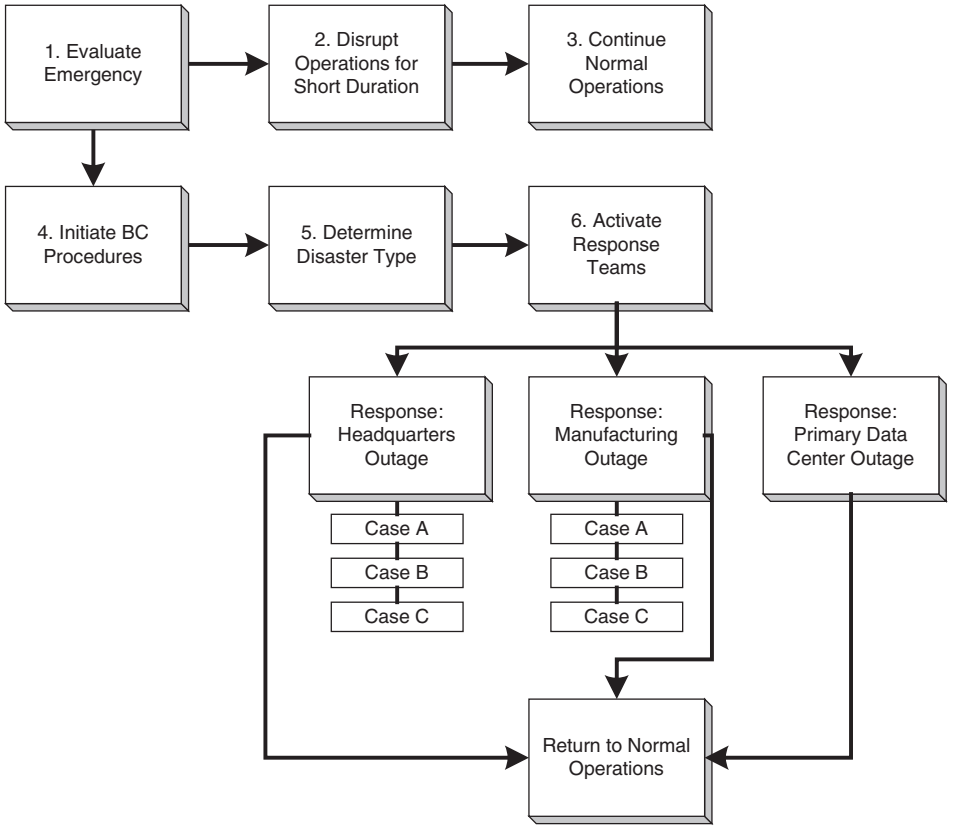


EXHIBIT 58.22 Recovery Flow

they reside. The diagram intends to clarify that categories of functions continue across time boundaries, and new categories are activated at their designated times.

In this example, Case A represents the time period from 1 hour after an event to 1/2 day after the event. In Case A, only Category 1 functions are activated. If the event lasts longer than 1/2 day, up to 2 days, then after the 1/2-day period, the organization enters Case B and activates Category 2 functions but leaves Category 1 functions active. If the event duration is greater than 2 days, then Case C shows that Category 1 and 2

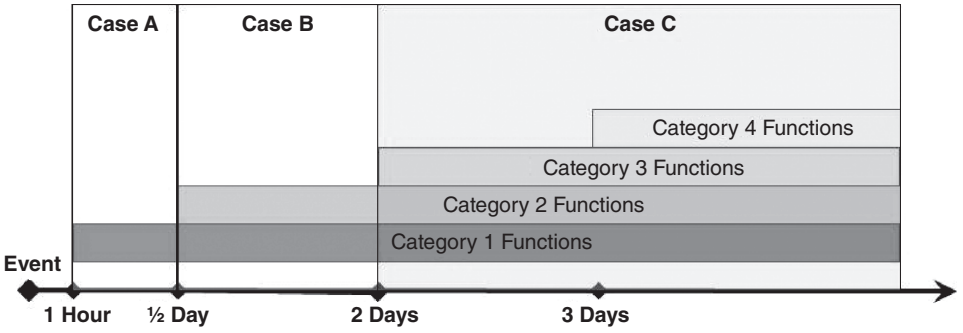


EXHIBIT 58.23 Cases and Categories

functions continue, and Category 3 and 4 functions are activated at the 2- and 3-day points.

58.7 CONCLUDING REMARKS. This chapter has introduced a methodology for analyzing mission-critical functions in an organization and for structuring business continuity plans that reflect organizational needs and resources. Chapter 59 in this *Handbook* continues the discussion with details of how to recover full functionality should a disaster occur.

58.8 FURTHER READING

Barnes, J. C. *Business Continuity Planning and HIPAA: Business Continuity Management in the Health Care Environment*. Rothstein Associates, 2004.

Barnes, J. C. *A Guide to Business Continuity Planning*. Wiley, 2001.

Butler, J. G., and P. Badura. *Contingency Planning and Disaster Recovery: Protecting Your Organization's Resources*. Computer Technology Research Corp., 1997.

Cimasi, J. L. *Disaster Recovery & Continuity of Business: A Project Management Guide and Workbook for Network Computing Environments*. CreateSpace Independent Publishing Platform, 2010.

Engemann, K. J., and D. M. Henderson. *Business Continuity and Risk Management: Essentials of Organizational Resilience*. Rothstein Associates, 2011.

Fulmer, K. L. *Business Continuity Planning: A Step-by-Step Guide with Planning Forms on CD-ROM*, 3rd ed. Rothstein Associates, 2004.

Hiles, A. *The Definitive Handbook of Business Continuity Management*, 3rd ed. Wiley, 2010.

Hiles, A. *Business Continuity: Best Practices—World-Class Business Continuity Management*, 2nd ed. Rothstein Associates, 2003.

Hiles, A., ed. *The Definitive Handbook of Business Continuity Management*, 2nd ed. Wiley, 2008.

Myers, K. N. *Manager's Guide to Contingency Planning for Disasters: Protecting Vital Facilities and Critical Operations*, 2nd ed. Wiley, 1999.

Roebuck, K. *Business Continuity and Disaster Recovery: High-Impact Technology—What You Need to Know: Definitions, Adoptions, Impact, Benefits, Maturity, Vendors*. Tebbo (2011).

Sayed, A., and A. Sayed. *Business Continuity Planning Methodology*. Sentryx, 2003.

Snedaker, S. *Business Continuity and Disaster Recovery Planning for IT Professionals*, 2nd ed. Syngress, 2013.

Watters, J. *The Business Continuity Management Desk Reference: Guide to Business Continuity Planning, Crisis Management and IT Disaster Recovery*. Leverage Publishing, 2010.

CHAPTER 59

DISASTER RECOVERY

Michael Miora

59.1 INTRODUCTION 59·1

59.2 IDENTIFYING THREATS AND DISASTER SCENARIOS 59·1

59.2.1 Threats 59·2
59.2.2 Disaster Recovery Scenarios 59·3

59.3 DEVELOPING RECOVERY STRATEGIES 59·6

59.3.1 Recovery Phases 59·7
59.3.2 Range of Strategies 59·9
59.3.3 Data Backup Scenarios and Their Meanings 59·13

59.4 DESIGNING RECOVERY TASKS 59·14

59.4.1 Beginning Sequence 59·14
59.4.2 Middle Sequence 59·16
59.4.3 End Sequence 59·18

59.5 IMPLEMENTATION AND READINESS 59·20

59.6 CONCLUDING REMARKS 59·21

59.7 FURTHER READING 59·22

59.1 INTRODUCTION. In Chapter 58 in this *Handbook*, the importance of a business impact analysis (BIA) and the method of preparing one were described. Once the preliminary groundwork is finished and the BIA analysis is complete, the next step is to design specific strategies for recovery and the tasks for applying those strategies. In this chapter, we discuss the specific strategies to recover the Category I functions, the most time-critical functions identified during the BIA, as well as the remaining lower-priority functions. We examine the traditional strategies of hot sites, warm sites, and cold sites as well as a more modern technique we call *reserve systems*. We describe how to make good use of Internet and client/server technologies and of high-speed connections for data backup, for making electronic journals, and for data vaulting. We develop the recovery tasks representing the specific activities that must take place to continue functioning and to resume full operations. These tasks begin with the realization that there is, or may be, a disaster in progress, continue through to full business resumption, and end with normalization, which is the return to normal operations. We examine a set of tasks taken from a real-world disaster recovery plan to illustrate how each task fits into an overall plan, accounting for anticipated contingencies while providing flexibility to handle unforeseen circumstances.

59.2 IDENTIFYING THREATS AND DISASTER SCENARIOS. Threat assessment is the foundation for discovery of threats and their possible levels of impact.

Threat assessments can vary from rigorous analyses of natural events and the probabilities associated with them to informal surveys of recent disasters and their regional characteristics. The rigorous analysis will yield detailed results but will require significant effort to perform the required research and to interpret the results. The informal survey can yield sufficient information for the disaster recovery planner, and it requires significantly less effort.

59.2.1 Threats. Local fire, police, and other emergency response services can provide the data needed for an informal survey. Flood district management agencies can provide data regarding floods of all types, whether caused by storms, draining, or other factors. Police departments can provide information about building or neighborhood closures due to man-made catastrophes, whereas city and county governments can provide information about emergency declarations over a period of years.

Compiling this list of threats accomplishes two objectives: It helps define the different levels of impact, and it helps to develop a risk mitigation plan. The mitigation techniques outlined in the plan are intended to lower the damage caused by the likeliest disasters. For example, an area prone to seasonal flooding can reduce the risk of computer outages by placing computer installations and important equipment above the flood level. Facilities located in areas prone to power failures can install UPS (uninterruptible power supply) equipment with higher capacities than usual for the type of equipment being protected.

Many threats can have important consequences to information-processing systems. Physical security risks stand out among them as the highest-profile disasters. These are also the most highly dependent on geography and include events such as fire, flood, and earthquake.

Exhibit 59.1 lists a large number of threats to help the disaster recovery planner begin the process of listing threats that apply to specific facilities.

EXHIBIT 59.1 List of Threats

Accidents	Denial-of-service attacks	Kidnapping/hostages
Acts of God	Disgruntled employees	Labor disputes
Aircraft accident	Drug abuse	Local area disaster
Alcohol abuse	Dust storms	Malicious damage
Area evacuation	Earthquakes	Mass illness
Arson	Elevator failure	Municipal service failure
Asbestos	Embezzlement	Nuclear accident
Boiler explosion	Environmental controls fail	Postal strike
Bomb threat	Evacuation	Power failure
Brownout	Falling object	Regional disaster
Building inaccessible	Fire	Sabotage
Chemical spill	Flood	Sanctions
Civil unrest riots	Fraud	Sandstorm
Cold weather	Hacker attack	Snowstorm
Communications failure	Hardware failures	Strike
Computer virus	Heat	Terrorism
Construction disturbance	Humidity	Tornado
Crime	Hurricane	Utility failure
Data diddling	Ice storms	Volcano
Delivery interruptions	Industrial espionage	Water damage
Demonstrations	Information compromise	Water supply failure

59.2.2 Disaster Recovery Scenarios. Each of the threats described in Exhibit 59.1 causes specific damage, with the scope of damage and damage characteristics well known and predefined. Once the characteristics of specific threats are known, then mitigation measures can be applied. One of the important mitigation techniques is the application of survivable technologies to mission-critical functions. Survivable technologies applied to the most critical of the functions (Category I) can prevent outages or at least significantly lower the probability and duration of the outage. For example, client/server and distributed architectures can be implemented with geographic data distribution, so that the loss of a single data center will not disrupt operations but may only degrade them. One such strategy is the geographically dispersed placement of two computers operating on a single main network with functionality shared between them. If backup data are regularly logged from one machine to another, then both systems maintain current databases. If one system is disrupted, then the other can assume full functions; the response times may degrade, but the functionality will not be lost. Replicated architectures with automatic failover are an even better alternative. The desirability of such measures was amply demonstrated by the catastrophe of September 11, 2001, again by Hurricane Katrina, and yet again in the Southern California fires of 2007.

Escalation scenarios are mechanisms that map the expected duration of failures against the requirements for operational continuity. With well-known threats as described, and the ensuing and defined outage durations that can be characterized, disaster declaration and escalation points can be calculated and presented. Timelines must be carefully constructed to leave little room for doubt.

Exhibit 59.2 provides an example of an escalation timeline. Note that the figure is based on data from Chapter 58, and from Exhibit 58.17 in particular. In this exhibit, the top, horizontal row of boxes corresponds to recovery time estimates. For each column, specific initiation parameters are established. The operators and engineers can, at a glance, determine whether this particular situation requires a disaster declaration. For example, an estimated downtime of less than one day requires no recovery, since in Exhibit 58.17 there are no functions requiring recovery in less than one day. The key factors are that the operations personnel must be trained in the use of the procedures and must have ready access to contact information for disaster recovery declaration authorities.

Disaster recovery scenarios and activities are dependent on the scope of the damage. Therefore, it is necessary to group specific events according to a classification of the damage they inflict on a facility. Once this is accomplished, the disaster recovery planner can identify declaration, escalation, and recovery activities based on the specific threats and their resulting disaster scenarios. Classifications of damage include the duration of an outage and the scope of its effects.

Exhibit 59.3 summarizes some threat durations and provides some sample classifications of damage. The first column of the exhibit lists a sampling of threats, and

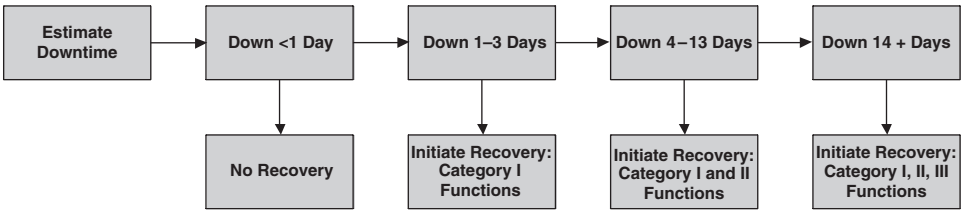


EXHIBIT 59.2 Escalation Timeline

EXHIBIT 59.3 Sample Classifications of Damage

Event Description	Possible Scope of Effect	Potential Duration of Effect
Aircraft accident	Building to local area	Hours to days
Asbestos	Partial to full building	Days to weeks
Bomb threat	Building to local area	Hours to days
Chemical spill	Building to lacal area	Hours to days
Civil unrest/riots	Local to regional	Days to weeks
Cold weather	Local to regional	Days
Communications failure	Building to regional	Hours to days
Computer virus	Systems	Hours to days
Data diddling	Systems	Hours to days
Earthquakes	Local to regional	Days to weeks
Environmental controls fail	Partial to full building	Hours to days
Fire	Partial to full building	Days to weeks
Flood	Local to regional area	Days to weeks
Hardware failures	Systems	Hours to days
Hurricane	Local to regional area	Days to weeks
Ice storms	Local to regional area	Days to weeks
Municipal service failure	Local	Hours to days
Nuclear accident	Regional area	Days to weeks
Snowstorm	Local to regional area	Days to weeks
Volcano	Regional area	Weeks
Water damage	Building to local area	Hours to days

the second column classifies each threat according to the possible scope of its effect. The set of possibilities contains systems only, partial or full building, local area, and regional area.

A systems-only disaster leaves the facility unaffected but partially or fully disables the system's operations. Breaches of logical security usually affect only systems, as do hardware failures. The exhibit does not attempt to differentiate between types of systems and number of systems affected, but the disaster recovery planner must do that classification for each facility and for the interactions of facilities. A characteristic of systems-only disasters is that sometimes they can travel across physical boundaries, such as when the incident involves a security breach, a connectivity disruption, or a software failure.

As an example of a system disaster, in 2007, one major company experienced just such a failure during the testing of a dry pipe sprinkler system. When the system was pressurized, the pipes burst. The data center was not serviced by the sprinkler system; it had its own system in place. However, one pipe crossed over the data center, a fact thereto unknown to the systems management personnel. When that pipe burst, the entire data center, which included a raised floor, was quickly flooded, making all the equipment wet from top and bottom. Fortunately, this company had prepared for the loss of any one of its data centers by building a load-balanced system, so that the loss of any one data center would cause only minimally reduced response times.

Partial- or full-building disasters consist of events that cause physical damage to a building. Fires usually affect only one building at a time. Even detection of asbestos, which most often occurs during building maintenance or improvement, is seldom a multibuilding affair. Local and regional area disasters are difficult to distinguish from

each other. Generally, a disaster that affects a single geographic area, such as a city or neighborhood, is a local area disaster, whereas a disaster affecting more than one geographic entity is a regional disaster. The importance of this demarcation is that local area disasters at a company site may affect some of the employees personally, although most will be unaffected, except as it relates to company business. A regional disaster usually affects most employees' homes and families as well as the organizations that employ them. A municipal services failure clearly affects a single geographic entity, and employees living outside of that municipality are unlikely to be affected. Conversely, a major storm resulting in floods, or a large earthquake, will affect residents and businesses over a wide area, probably leaving employees at work concerned with the health and safety of their families and homes. The availability of employees to handle a disaster is much greater when the disaster is local, or of smaller scope, than during a regional disaster.

The third column in Exhibit 59.3 estimates the duration of the effect, ranging from hours to days to weeks. These are inexact and somewhat vague guidelines only. Although fire is listed as having a duration ranging from days to weeks, a very small fire may be managed and the building reoccupied in only a few hours. Similarly, an earthquake is listed as having a duration of days to weeks, yet a major earthquake may have large impacts lasting many months. After the 1994 Northridge earthquake in Southern California, for example, major highway arteries were closed for many months, requiring employees to telecommute or spend hours per day traveling. This affected the operations of many companies, forcing them to redesign their workflows. The Seattle earthquake of February 2001, however, had no such long-term effects. Many companies initiated disaster recovery procedures to handle the commuting problem, even when there had been no physical damage to the corporate facilities. The terrorist destruction of the World Trade Center twin towers in New York on September 11, 2001, is perhaps the most obvious example of how the destruction or closure of major roads and other transportation means affected operations of companies not otherwise affected by the terrorist act itself.

The disaster recovery planner must analyze enough specific threats against the corporate operations and physical plant to build a series of disaster scenarios, which become the guiding forces of the recovery activity definitions. Some planners design disaster recovery scenarios to suit each disaster threat. Floods call for one set of activities, fires another. This can result in an extensive set of scenarios, each requiring training, testing, and maintenance. This technique may be too complex. Another similar approach uses active scenarios chosen from a set of only five possible levels of disasters: systems only, partial building, full building, local, and regional. Each organization should try to reduce the number of disaster scenarios, generally to no more than three, and many can handle emergencies adequately using only two scenarios. For most organizations, a systems-only disaster is not restricted to computer systems, but may include phone systems and telecommunications systems. Systems-only cases must be defined in the context of the organization's unique structure.

For example, a manufacturing facility may require three scenarios: systems only, partial building, and full building. Perhaps for the manufacturer, a loss of systems can be handled with a hot site for computers and a fast, high-bandwidth connection to the shop floor. A partial building disaster, affecting only office space, can be handled similarly, but with additional office space for support functions. A full-building disaster that damages or destroys manufacturing capability requires a completely different set of continuation and recovery procedures. As another example, a corporate office facility may require just two scenarios—systems only and building disaster—where the

building disaster is a combined building, local, and regional scenario. The systems hot site can handle the systems-only disaster. The building disaster scenario will require that all functions be restored at a hot site that provides full office functions from phone to fax to desks and copiers in addition to computer systems and networks.

The planner defines the sets of disasters and develops recovery scenarios for those sets. This defines the number of scenarios required; then the disaster recovery activities can be designed.

59.3 DEVELOPING RECOVERY STRATEGIES. The BIA defines the timeline for recovery for each category of functions. The recovery strategies depend on factors such as: the complexity of the functions, the amount of information required to carry them out, the number of people involved in performing the function, and the amount of interaction between this function and other functions. The strategy of choice for each function is based on the timeline and the specific factors relating to it. The tasks required to activate the function immediately after a disaster follow directly from the BIA timeline and from the various factors that have been described.

Recovery strategies are defined differently from continuation strategies. The recovery strategy is the overall plan for resuming a function, or set of functions, in a near-normal work mode. The continuation strategy is the plan for immediate or nearly immediate operations of a key function, even if it results in significantly degraded performance or limited capability. The continuation strategy is most often applied to a small set of functions rather than to a more complete set, and it is often only temporary. Although the recovery strategy may be used for several weeks or longer, the continuation strategy typically survives for hours or days.

Traditional recovery strategies include hot sites, warm sites, and cold sites. The most common recovery strategy uses a commercial service provider to maintain an on-demand operating capability for well-defined systems. There are a variety of such companies, some of which provide a standby mainframe or midrange system, along with selected peripherals and telecommunications facilities on a first-come, first-served basis. Others guarantee availability by providing a dedicated facility.

The operating capability required is defined at the outset. Processor type is specified down to the model level. Required disk space, hard drives, tape drives, ports, connectivity, and all other required elements must be specified explicitly. Timeline requirements must be stated precisely, with rehearsal, practice, and drill times negotiated in advance. Geographic location of the backup systems is often determined at the time of disaster declaration, and connectivity is usually provided to corporate fixed facilities. The obvious advantage of this strategy is ready access to maintained systems and environments. The principal disadvantage is cost.

A continuation strategy usually depends on internal resources rather than on commercially available providers. Continuation strategy requires that equipment be available instantly, with current data. "*Instant availability*" means that the systems must be local, or at least connected, yet survivable. The continuation strategy is becoming more common as microcomputer-based systems assume greater roles in corporate computing structures. The continuation strategy often relies on reserve systems, a concept pioneered over the past five years. This strategy, described in Section 59.3.2.8, uses high-performance microcomputers residing off site at homes, branch offices, and other locations easily accessible to employees.

The reserve system provides immediate, but limited, functionality so that operations can be continued after a disaster while awaiting full restoration or recovery.

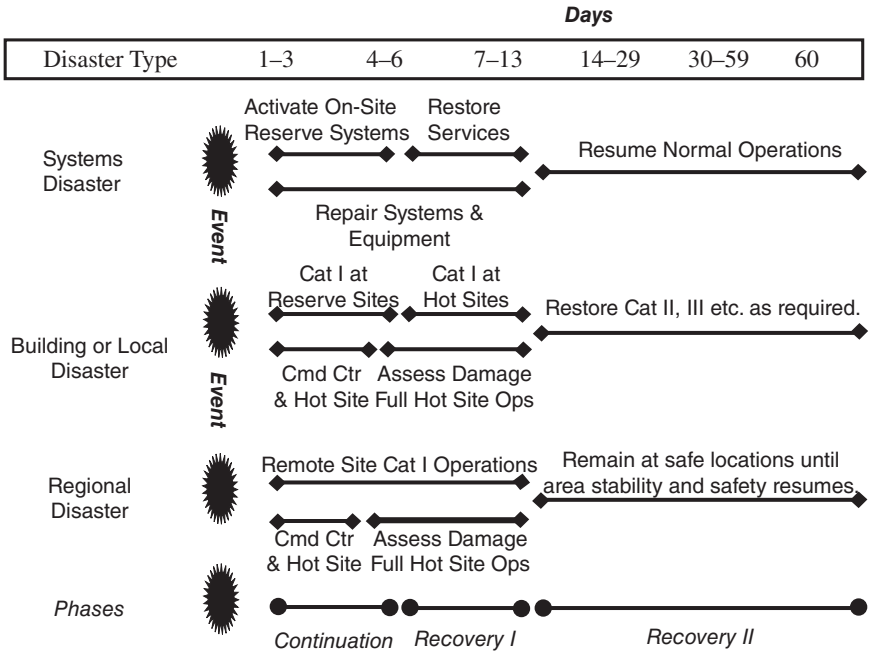


Exhibit 59.4 Recovery Phases

59.3.1 Recovery Phases. Orderly recovery requires organized processes and procedures. The BIA uncovers the time sensitivity and criticality of functions, so that continuation and recovery strategies can be designed. Activation of these strategies is performed systematically in time-phased increments.

Exhibit 59.4 shows an example of such phasing. The figure is based on Exhibit 58.17, and illustrates a common schema that consists of three scenarios. Each scenario represents a collection of disasters according to their effects on the company. The first, a systems disaster, represents an effect that has interrupted computer and related systems but has left basic services such as heating and air conditioning, water, power, and life safety intact. The second scenario is the building or local disaster. In this case, a major portion of the building, the entire building, or even several surrounding buildings are rendered unfit for occupancy. This may be due to a problem with the building or with the local area, including fire, toxic spills, minor flooding of surrounding areas, terrorist action, and other such events. The third scenario is the regional disaster, in which some calamitous event affects a wider geographic area. Examples of such catastrophes are earthquakes, hurricanes, floods, large fires or firestorms, and heavy ice and snowstorms.

Each of these scenarios requires different recovery strategies and tasks. However, the overall recovery phases remain similar. The exhibit shows that there are three basic phases: continuation of critical (Category I) functions, recovery of critical functions, and recovery of other functions. These are labeled near the bottom of the figure as "Continuation," "Recovery I," and "Recovery II."

The Continuation phase begins at the time of the disaster. For each scenario, the goal of the continuation phase is to support the Category I functions as best as possible and within the time frames defined by the BIA. In this example, the BIA timeline for Category I functions is one to three days. For the systems disaster and the building or

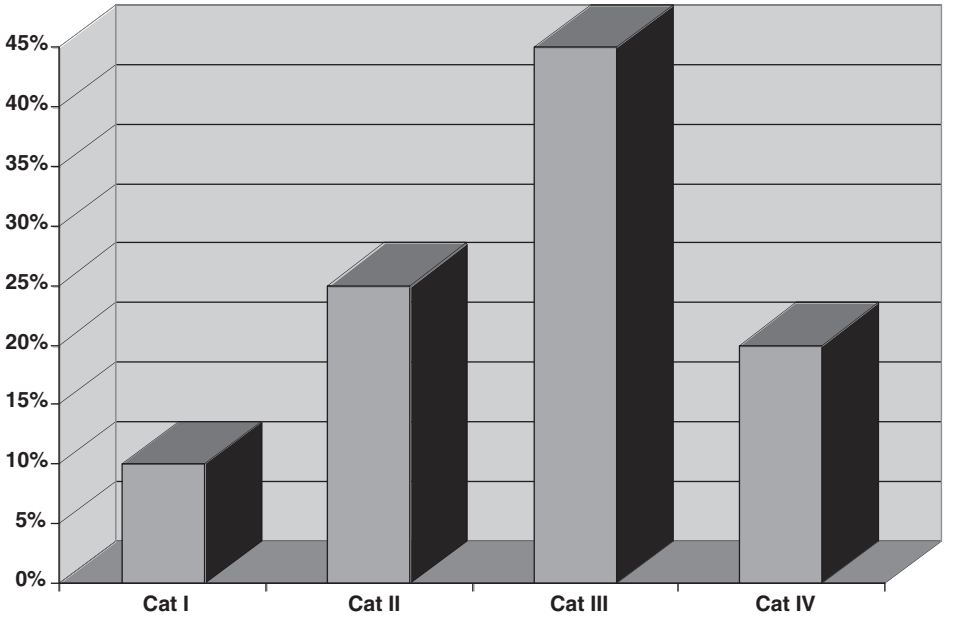


EXHIBIT 59.5 Spread of Functions across Categories

local disaster scenarios, the Continuation phase depends on the local or onsite reserve systems. For the regional disaster, where local functioning may be impossible, the Category I functions must be supported by remote sites.

The recovery phase is further divided into two subphases. Recovery I phase restores full or nearly full functionality to the Category I functions that were partially supported by the reserve systems. In the example, the Recovery I phase begins after all continuation functions have been stabilized. In practice, it is more common for the Recovery I phase to begin shortly after the disaster event in order to minimize the duration during which the Category I functions must operate in a degraded fashion. The Recovery II phase is used to recover other functions in the order they are needed, as defined by the BIA. Exhibit 59.5 shows the typical spread of functions across categories. Typically, Category I functions constitute approximately 10 to 15 percent of the functions performed in a facility or campus. The Category II functions consist of functions that are very important and time critical but can be delayed for a short time. These functions significantly outnumber the Category I functions, sometimes reaching 20 to 35 percent of facility functions. Category III and, if applicable, Category IV functions are much greater in number, often comprising more than half of the facility functions performed.

The BIA process assessed the cost contribution of each category of functions to the total losses the company would sustain after a disaster without a plan. Exhibit 59.6 compares the number of functions per category to the total impact of the functions in each category on operations in a typical organization. Although Category I functions number the fewest of all categories, their impact is the greatest. The Category III functions typically rank second in contributions to losses, with Category II third in rank. The critical functions that constitute Category I are ranked first precisely because of their high impact and time criticality. Category III is typically the largest group, so that even though the impact of each function may be small, the large number of

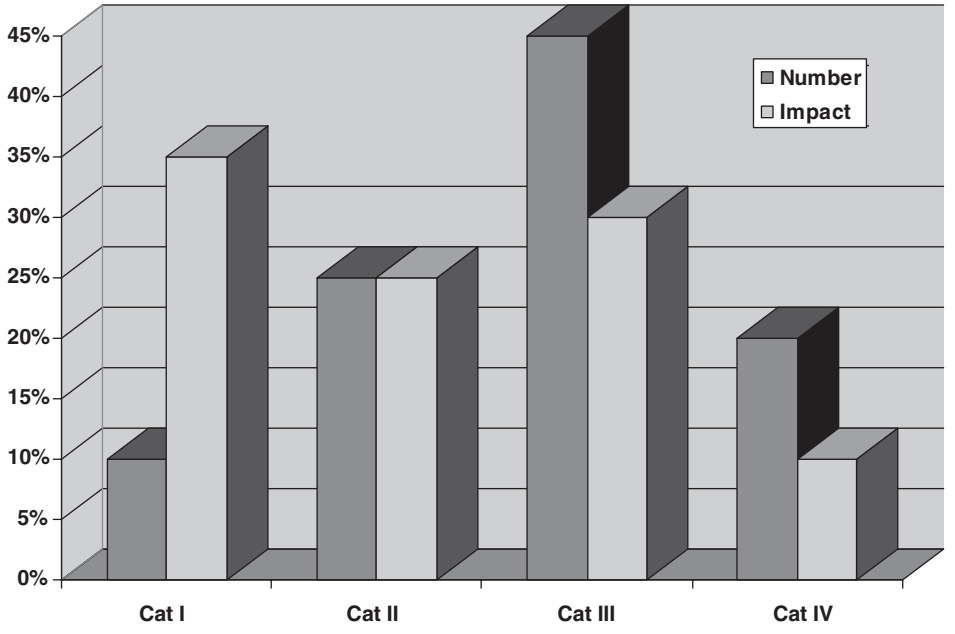


EXHIBIT 59.6 Number versus Impact of Functions

functions increases the total effect of that group. Category IV tends to be a fairly small set of functions, so typically it is small in number and last in impact.

In the example, the systems disaster phases reflect certain assumptions. The Continuation phase begins at the time of the disaster and continues for up to three days. The Recovery phase is scheduled to begin no later than the fourth day, with full recovery by the seventh day, at which time a return to normal operations is expected. The underlying assumption is that the off-site reserve systems have been replicated on site. This is shown by the caption "Activate On-Site Reserve Systems" in Exhibit 59.4. An additional assumption in this example is that the system can be recovered in less than seven days. Some analysis would have been performed to show that new equipment or equipment repairs could be obtained fairly early during the first seven days, leaving enough time for service restoration, including data recovery, to be completed by the seventh day. In this example, it is also clear that the Category II functions are required beginning on the seventh day. This is illustrated by the building or system disaster Recovery II phase, which begins on the seventh day and is labeled as restoring Category II, III, and other functions.

59.3.2 Range of Strategies. Three general areas of traditional recovery strategies address the needs of mainframes, midrange systems, and PC LAN systems. There is a rich history and many years of experience in providing recovery for large, mainframe systems, with strategies that include hot sites, cold sites, reciprocal or mutual aid agreements, and commercial recovery facilities, as will be described.

Strategies for recovery of midrange systems, including distributed systems, minicomputers, and workstations, include all the same strategies as those used for mainframe systems, along with mobile data centers, internal redundancy, failover arrangements, and agreements with manufacturers for priority shipping of replacement equipment.

PC LAN system recovery strategies are the newest of the three areas. They typically make use of combinations of capabilities, including quick replacement, off-site equipment storage (equipment sparing), and hot sites.

59.3.2.1 Cold Sites. A cold site is a room or set of rooms in a *ready-conditioned* environment. It is a fully functional computer room with all of the required facilities, including electrical power, air conditioning, raised floors, and telecommunications, all in operating condition and ready for occupancy. Missing from this fully functional room, because of high costs, is the computer equipment, including the processors and all required peripherals. Cold sites can be owned by one company or shared by several companies.

The major advantage of a dedicated cold site, somewhat diminished in the shared case, is simply that for a relatively low acquisition or leasing cost, the site is guaranteed to be available over the long term to the owner or lessee of the site. To be effective, a cold site must be distant enough from the main facility so that a disaster that makes the primary facility unusable will likely not affect the cold site.

There are several disadvantages to cold sites. First and foremost is that ordering, receiving, installing, and powering up the computer system can take many days. Once the system is functional, the cold site becomes a dedicated facility ready to perform all functions in a controlled environment, but the time required to achieve this state can stretch into a week or more. Few organizations can rely on a cold site as their primary vehicle for disaster recovery. A secondary disadvantage of cold sites is the inherent inability to test the recovery plan. An untested plan is an unreliable plan—testing a cold site would require obtaining all of the requisite equipment, installing the equipment, and performing operations on the new equipment. Few organizations can afford the costs associated with such testing. There are also hidden pitfalls with this strategy. Key equipment may be unavailable for immediate delivery. Communication lines may be untested and unreliable. This strategy is generally not a desirable first line of defense. It can, however, be a part of a larger overall strategy that includes other types of sites.

59.3.2.2 Hot Sites. A hot site is a facility ready to assume processing responsibility immediately. The term usually refers to a site that contains equipment ready to assume all hardware functions, but requiring massive restoration of data and an influx of personnel to operate the equipment. The hot site cannot, therefore, begin total processing instantaneously. In actuality, the site is only warm, with an ability to get hot fast. A hot site can be dedicated to one organization, shared by several organizations, or leased from a specialty company, called a *commercial recovery service provider*. The provider option is described in a later section.

The primary advantage of a hot site is the speed of recovery. The time to resume processing is defined as the time to reach the facility with people and backup media, plus the time to restore data and programs from the backup media, to test the operations, and to go live with a fully functional system. For larger systems, this period can range from less than two days to almost a week.

The primary disadvantage of a hot site is the cost associated with acquiring and maintaining the fully equipped site. For this reason, most organizations choose to share the cost of a hot site with other organizations, which may be sister companies or only neighbors. In any event, there is a large cost associated with maintaining the site and ensuring that it is updated with every change to each of the participant's requirements, while still maintaining compatibility with all of them. One of the most common solutions to this problem is to use a service provider.

59.3.2.3 Reciprocal Agreements. Reciprocal agreements were often used in the earlier decades of disaster recovery planning but are uncommon today. A reciprocal agreement is an arrangement between two or more companies in which each agrees to make excess capacity available to the others in case of a disaster. The major advantage is the apparent low cost of this solution. The major disadvantage is that these arrangements rarely provide the needed computing power. A major issue with these arrangements is maintaining compatible systems. If one company changes a processor, it may find that its partners' systems cannot perform adequately, even in degraded mode. If the host company faces a crisis or deadline of its own, a reciprocal company may find itself with no computer power at all. These arrangements are seldom testable, because it is the rare company willing to shut down operations to help a partner perform a disaster recovery test.

59.3.2.4 Internal Redundancy. A strategy of internal redundancy requires that a business have multiple facilities, geographically dispersed, with similar equipment in each site. If there are data centers at several sites, then the alternate data centers may be designed with excess capacity to support a potential failure at another site.

The major advantage of internal redundancy is that the organization maintains complete control of all equipment and data, without relying on any outside company to come to its aid. The excess capacity at the various alternate sites must be carefully protected, and management must exercise diligence in budgeting and operations. Careful intracompany agreements must be crafted to ensure that all parties are aware of, and agree to, the backup arrangements. Internal redundancy can be an effective solution in cases where temporarily degraded processing still can provide sufficient support to meet timeline requirements. If degraded performance is not an acceptable option, then the cost of the excess capacity will probably be too high. If reasonable degradation is acceptable, then those costs can be manageable.

Internal redundancy can also be difficult to test. Testing requires that processing be shifted to a recovery mode. Unlike external, separate computers, all of these redundant systems would be operational. Testing one disaster recovery plan requires affecting a minimum of two corporate locations. A failed test that causes a system crash or other problem can have damaging consequences.

Internal redundancy can overcome its difficulties if the redundancy is part of a load-balanced system with automatic or semiautomatic failover. In such cases, normal processing is already split between the primary and backup system so that both systems are fully operational at all times. The recovery requirements add an additional dimension to the nominal load-balancing configuration. Namely, there can be no system or system component that is unique to only one site. The most common load-balancing implementations may still rely, for example, on a database housed at only one site. This is insufficient for recovery purposes. If full duplication can be achieved, then the sudden loss of one of the systems causes degradation rather than loss of functionality. There are, however, often technical issues that many organizations must overcome prior to the implementation of such load balancing with full redundancy. For example, some legacy systems do not support multiple points of processing; some modern implementations may require a single database instead of multiple, overlapping databases, as is required for full load balancing with redundancy.

59.3.2.5 Mobile Data Centers. Mobile data centers are transportable units such as trailers outfitted with replacement equipment, air conditioning, electrical

connections, and all other computer requirements. Mobile data centers are used most often for recovery of midrange and PC LAN systems. The primary advantage of the mobile data center is that it can be activated quickly at reasonably low cost. The primary disadvantages are the expense of testing such a facility and the possibility that a local or regional disaster will prevent successful activation. Deploying a mobile data center requires careful planning. Land must be available to accommodate the transportable units, with outside parking lots as the most common resource. Local government and municipal regulations must be researched in advance to ensure that such units do not violate ordinances, and can arrive as certified for immediate occupancy. External power and communications hookups also must be available.

59.3.2.6 Priority Replacement Agreements. Some computer vendors support priority equipment replacement agreements. These are arrangements in which the vendor promises to ship replacement equipment on a priority basis. For midrange systems, this is often an agreement to send the next-off-the-line system to the priority customer. The major advantage of this strategy is its low cost. However, if the vendor is not currently manufacturing the required system, or if the assembly line is down for any reason, and if equipment stocks are depleted, there may still be a significant delay in receiving the equipment. This is the major disadvantage. This strategy also assumes the disaster recovery plan makes an alternate facility available in case the primary facility is damaged along with the equipment being replaced.

59.3.2.7 Commercial Recovery Services. Commercial recovery service providers can support a combination of the strategies just discussed. These companies generally provide three major benefits: cost sharing, reduced management needs, and diminished risk of obsolescence. First, they spread facility costs across multiple subscribers so that each subscriber saves as compared with building and maintaining a comparable, privately held capability. Because all subscribers share the same physical space, each pays less than would be needed to maintain such a site independently.

The second major benefit of using a commercial provider is that these companies eliminate the need for the subscriber to manage backup resources. Management and maintenance of such a site by an individual business could be a heavy burden, but the provider's primary focus is on managing, maintaining, and upgrading the equipment and sites. The subscriber company can be assured that the equipment is exercised and serviced regularly and that peripheral equipment, power systems, and facility support structures are properly maintained. A properly run site will also provide security, safety, compliance with evolving rules and regulations, and competent staffing. The provider assumes full responsibility for these functions during normal operations and continues support during times of crisis. The subscriber brings its technical personnel, while the provider leaves its facilities staff in place.

The third major benefit centers around today's fast pace of hardware evolution. A subscriber company typically leases a hot site or other service for a five-year period. During that time, hardware platforms will evolve. The subscribing company can protect its lease investment by ensuring that system upgrades are reflected in the leased equipment configuration for reasonable extra charges. A business that provides its own hot site must upgrade the hot site whenever hardware, and sometimes software, changes are made to the operational systems.

The disadvantage of commercial recovery services is in the obvious risk that the hot site may not be available in an emergency. Indeed, if a local or regional disaster affects numerous subscribers, there could be significant contention for the provider's

resources. To address this issue, providers typically maintain hot sites in geographically dispersed areas. Although it is likely that in the case of a local or regional disaster a subscriber would need to use a hot site farther away than planned, it is unlikely that the subscriber would be left completely without the prearranged resources.

59.3.2.8 Reserve Systems. The newest of strategies is the reserve system, which is a small replica of a portion of an operational system meant for use during the first few days following a disaster. The reserve system provides continuation of key functions for short durations, although in degraded mode. The reserve system usually resides off site at an employee's home or at another corporate office. Another version of the reserve system is also kept on site. A reserve site may be equipped with a microcomputer or a minicomputer ready to assume functioning in case the primary system becomes unavailable. This meets the important criteria of a reserve system, which must be fast and easy to activate, simple to move, low in cost, testable, available, and highly reliable.

The reserve system concept was not feasible until client/server technology emerged and telecommuting with high-speed communications became accepted and readily available. Many large organizations provide employees the option of telecommuting on a regular basis so that they work at home sometimes and in the office at other times. Web-centric processing and remote application server technologies can provide powerful reserve systems for disaster recovery.

Proper security precautions must be taken to protect proprietary, confidential, and critical information stored on these reserve systems. Strong encryption protects against theft, while redundant systems protect against other losses. The reserve system is a quick-response, short-term solution intended to solve the problem of immediate continuation, even in the case where employees may be unable to travel outside their immediate residence areas. See Exhibit 59.7 for an overview of different recovery strategies.

59.3.3 Data Backup Scenarios and Their Meanings. Data backup is a key function in all system installations. The best recovery strategy, chosen to meet recovery timelines according to the BIA, is useless without a backup from which to restore and resume operations. Data backup is perhaps the single most critical element of a disaster recovery plan, yet only 31 percent of U.S. companies have backup plans and equipment. See Chapter 41 in this *Handbook* for details of backup scenarios.

EXHIBIT 59.7 Strategy Overview

Strategy	Activation	Cost	Testability	Availability	Reliability
Ideal strategy	Fast	Low	Excellent	High	Excellent
Reserve systems	Fast	Low	Excellent	High	Excellent
Internal redundancy	Fast	Medium	Poor	Medium	Good
Commercial providers	Fast	Varies	Excellent	High	Excellent
Hot site	Fast	High	Excellent	High	Excellent
Mobile data centers	Medium	Low	Medium	High	Good
Reciprocal agreements	Slow	Low	Poor	Low	Poor
Priority replacement	Slow	Low	Poor	Low	Poor
Cold site	Slow	High	Poor	High	Poor

59.4 DESIGNING RECOVERY TASKS. The disaster recovery plan becomes an operating document when the strategies and reserve capabilities are translated into specific actions to be performed during the disaster recovery process. These are the actions required to protect assets, mitigate damage, continue partial operations, and resume full operations. Most tasks begin with actions to evaluate the disaster and then to initiate recovery procedures. The example recovery task flow (Exhibit 59.8) lists a series of steps to evaluate the situation. These steps, numbered 1 through 5, determine whether to declare a disaster or not. The second set of steps, numbered 6 through 10, is used to initiate recovery procedures. These first two sets of steps are called the *beginning sequence*.

The beginning sequence leads to full activation of a disaster recovery effort. The next set of activities, called the *middle sequence*, institutes the actions necessary to perform business continuation and recovery. Usually multiple choices for the middle sequence correspond to the disaster scenarios developed. Each path in the middle sequence is executed for a particular scenario. At the end of the middle sequence, full operations are resumed as specified in the disaster recovery plan. In the exhibit, the middle sequence includes two distinct paths. The first path includes steps numbered 11 through 20, and the second path includes steps numbered 21 through 30.

The end sequence begins with damage assessments and continues through resumption of normal operations in the original facility, a semipermanent facility, or a new facility. The end sequence may begin while the middle sequence is still unfolding. The timing is dependent on the requirements specific in the BIA. The end sequence consists of two sets of activities performed in sequence. The first set consists of steps 31 through 38, and the second set consists of steps 39 through 59.

59.4.1 Beginning Sequence. The beginning sequence (Exhibit 59.9) helps the disaster recovery team in its first two major actions: evaluation of the emergency and initiation of the recovery procedures. Box 1, labeled "Disaster Evaluation," begins the process with a set of actions whose end goal determines what actions to take to meet the upcoming emergency. In step 2, the disaster recovery team determines which course of action to take: continue normal operations, disrupt operations for a short while and then resume normal operations, or initiate a full- or partial-scale disaster recovery. The decision is based on a set of predetermined criteria.

Box 2 is labeled "Evaluate Emergency." This is the step in which the disaster recovery team assesses the expected duration and extent of the upcoming or existing outage. Using a set of criteria developed as part of the BIA, and based on the assessment of the outage, the team determines which course of action to take. For example, if the emergency is a power outage, the disaster recovery team will likely call the power utility to determine the cause. If the outage is expected to last less time than the shortest recovery requirement, then it may make the most sense simply to disrupt operations for a short time without initiating recovery procedures. The comprehensive disaster recovery procedures provide sufficient guidance in a simple form to support the team's decision-making process while leaving room for real-time evaluation and decision making.

If the decision is to initiate disaster recovery procedures, then the flow proceeds through box 5, "Initiate DR Procedure," to the initiation phase represented by the actions in boxes 6 through 9. In step 6, the disaster recovery team determines which of the disaster scenarios applies; the team formally declares that type of disaster in step 7, and, in step 8, notifies the disaster recovery team and others, as required, of the

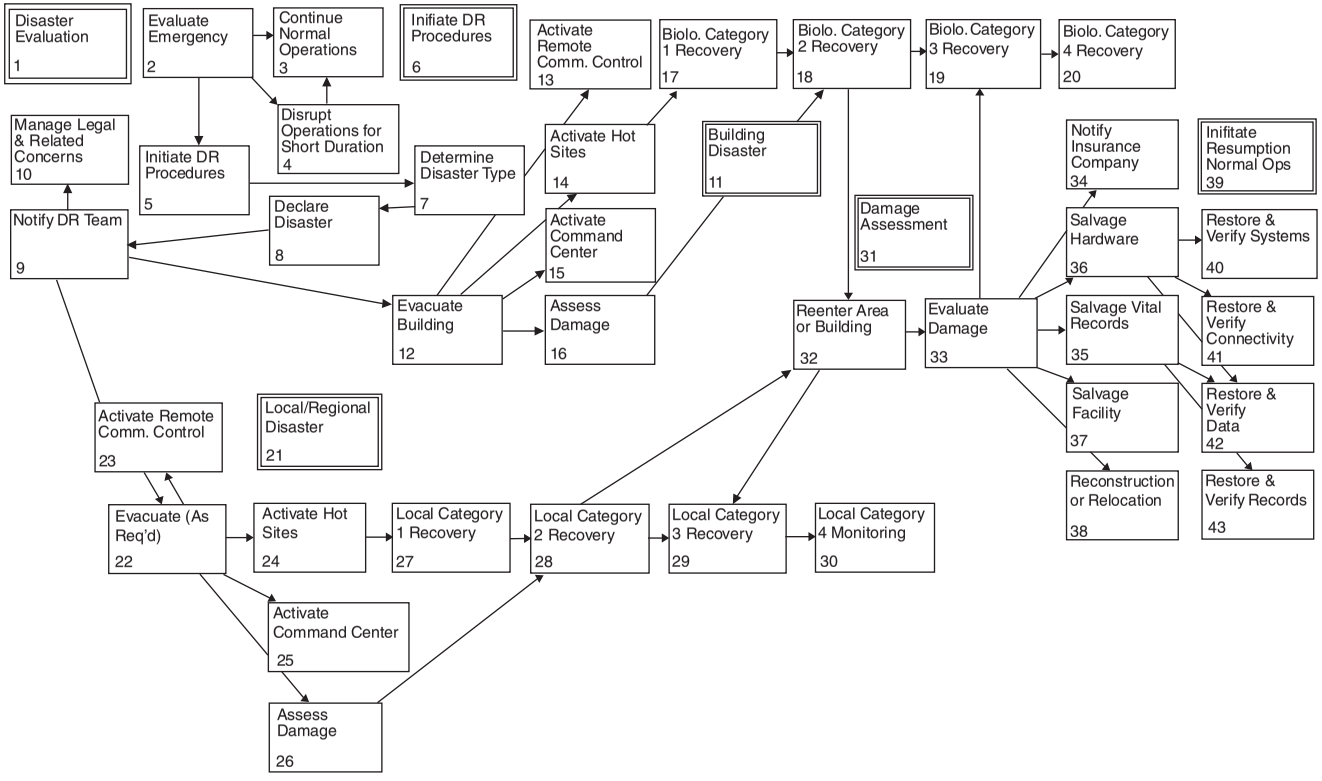


Exhibit 59.8 Recovery Task Flow

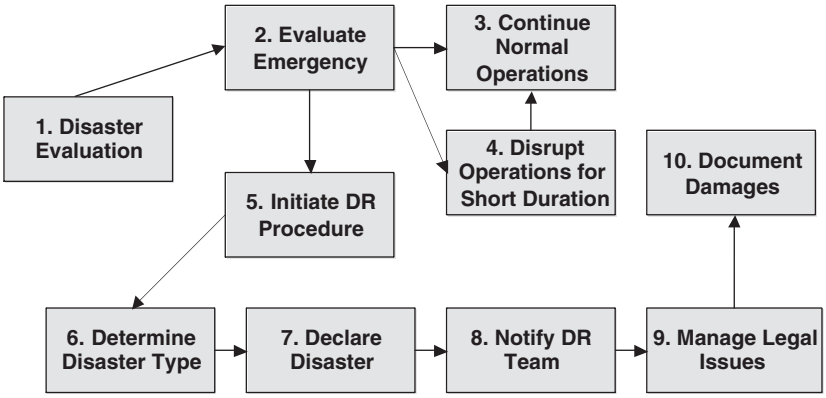


Exhibit 59.9 Beginning Sequence

decisions made and the next steps to take. All members of the team are pretrained and equipped with written procedures that anticipate each of these potential decisions. The only remaining step in the beginning sequence is box 9, "Manage Legal Issues." From the moment a disaster situation is considered, the team must document its decisions. In some industries, this is required by regulatory agencies; in public companies, it may be required for reporting to a board of directors' committee meeting. For all companies, it is important to maintain a clear record of the diligence used in planning and implementing recovery.

In addition to legal concerns, there are insurance concerns. Typically, money spent during the recovery process is reimbursable by insurance policies. This includes out-of-pocket expenditures by staff; the cost of travel, food, and lodging vouchers; and other allowed expenditures for meeting disaster recovery needs. In its haste to restore damaged systems, the team may neglect to document specific damages; once repaired, it is difficult to submit an insurance claim unless the damage is documented properly. Box 10 represents a step that begins early and continues throughout the disaster recovery process.

59.4.2 Middle Sequence. The middle sequence represents the activities that begin once the disaster recovery procedures are under way. In our example, there are two possible paths for the middle sequence to follow. The steps are labeled similarly, but the specific actions taken are quite different. There is an advantage to using a schematic representation of the steps and their relationships: The overall structure or shape can show at a glance the complexity or simplicity of the plan. It is best to be as uncomplicated as possible, even if the form and structure require larger content. Exhibit 59.10 describes the middle sequence, "Building Scenario."

The first step listed in the middle sequence, "Building Scenario," is to evacuate the building. Although this is clearly performed as early as possible, perhaps even early in the beginning sequence, it is presented here for completeness. Human safety is the primary concern of any recovery plan—no employee should ever be allowed to delay evacuation or reenter a building, except in the case of a qualified individual who is helping others evacuate or making a rescue attempt. The next four activities occur in parallel. They are represented by boxes 3 through 6, and include activation functions for remote communications, the hot sites, and the command center. These four steps also include a first effort at damage assessment.

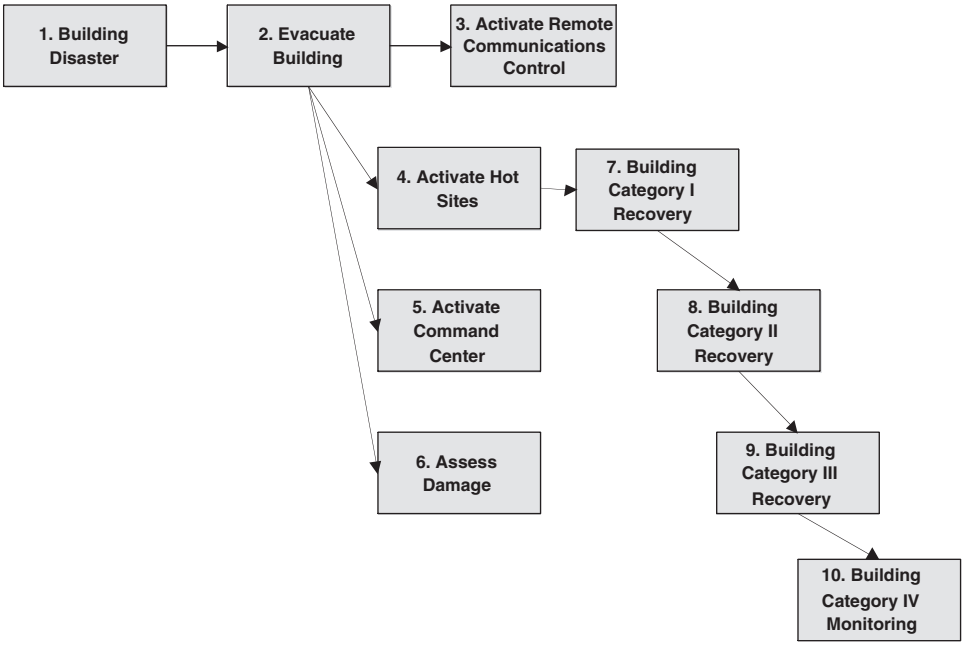


Exhibit 59.10 Middle Sequence, "Building Scenario"

Box 3 is labeled "Activate Remote Communications Control." This is the activity that reroutes digital and analog communications from their normal patterns to their emergency routes. This may include telecommunications lines, call center toll-free lines, normal business lines, data lines, and other communications media. Typically, arrangements are made in advance with the various telecommunications carriers so that a series of short phone calls can initiate the rerouting. If arrangements were not made in advance, then the rerouting will not meet timeline requirements and may significantly impair the efficiency of the disaster recovery team members.

Box 4, "Activate Hot Sites," represents the sets of actions required to inform the hot site provider that a disaster is in progress. Hot site activation must be performed at the earliest moment, so that the hot site can be available as soon as possible and so that, in cases of contention, the preferred hot site will be available.

Hot site activation usually obliges the company to pay certain activation costs and occasionally a minimum usage cost. If the disaster requires the hot site, then insurance often reimburses the company for related expenses, but if the declaration was erroneous, the company will suffer a financial loss.

The major hot site providers' services are quite reliable. There are few instances of hot site providers failing to meet their obligations, and many stories of near-heroic feats performed to meet customer needs. One provider even abandoned its own operations so that a client suffering from a regional disaster could occupy the provider's corporate space to continue its own operations. Hot site contracts usually include a proviso that in cases of contention, where more than one company is seeking the same disaster recovery space, the provider can send the late declarer to another, equivalent site. This means that the disaster recovery plan must include instructions for multiple hot sites. All maps, directions, instructions, and delivery requirements must be consistent with the assigned hot site location. Clearly, it is in the company's best interest to declare

early so that the hot site of choice is obtained. It would be prudent to err on the side of early declaration rather than late declaration, especially in cases of regional disasters.

Boxes 7 through 10 represent recovery of the various categories of functions. It is seldom possible or desirable to recover all functions simultaneously. The BIA determined and documented the timelines for all functions needing recovery. The procedures must reflect that analysis so that functions are in fact recovered *just in time*. These four boxes are labeled "Building Category Recovery" I to IV. Each box represents the full set of procedures required to bring all the functions in each category into operational mode. For example, if there are seven functions in the Category I group, then there will be seven sets of procedures required to implement Box 7. Category I functions usually require special efforts, such as reserve systems. Therefore, each Category I function may require two sets of activation procedures: the first for the reserve system and the second for the hot site.

At the conclusion of the middle sequence, all required functions would have been recovered. It may be that during the beginning sequence a determination was made that the disaster would require only recovery of a subset of functions. A small-scale power failure, for example, may require only Category I functions to be recovered. A weather disaster with no lasting effects may require Category I and Category II functions. This decision is made during the early phases of the recovery, and reevaluated continually throughout the recovery process.

59.4.3 End Sequence. The end sequence represented in Exhibit 59.11 is the long path back to normal operations. The end sequence consists of two major phases. The "Damage Assessment" set of activities comprises the first phase and includes boxes 2 through 8. The "Initiate Resumption of Normal Operations" set of activities constitutes the second phase and includes boxes 10 through 13. Although there is an

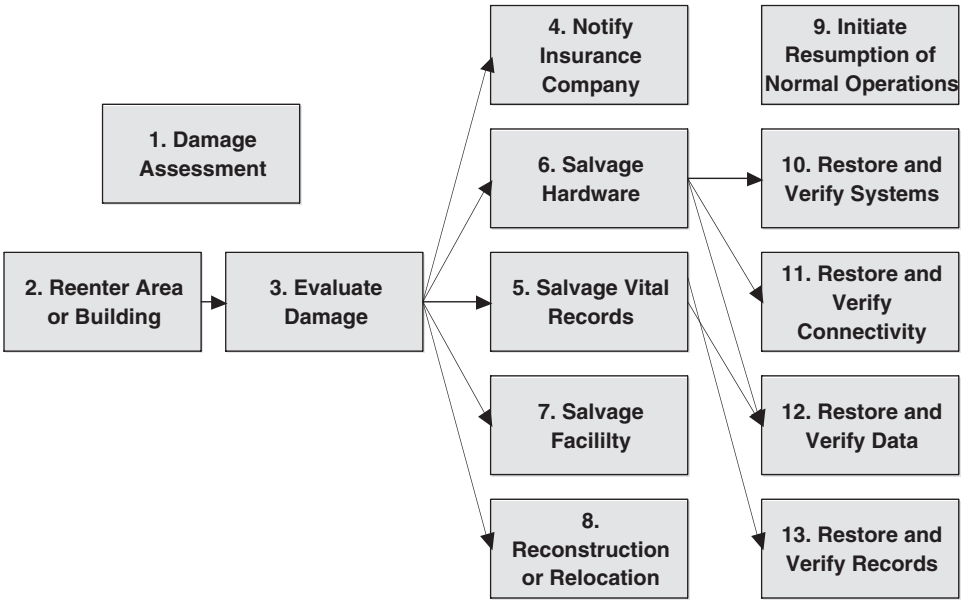


EXHIBIT 59.11 End Sequence

assess damage activity at the beginning of the middle sequence, that activity is a quick assessment of damage for the purpose of determining what recovery options need to be activated. This damage assessment activity is a comprehensive evaluation of damage to equipment, facilities, and records.

Box 2 is labeled "Reenter Area or Building." This activity needs supporting procedures to guide disaster recovery team members on what regulatory, municipal, state, and federal authorities control access into damaged facilities. Often access is restricted for long periods of time even when there is little or no apparent damage. For example, following an earthquake, structural engineers may tag a building as uninhabitable pending further analysis. As another example, after a fire, the fire marshal may prohibit entry pending analysis by chemical engineers to determine whether the fire produced hazardous materials. Fires often cause chemical reactions that transform normal, benign office materials into toxic chemicals. The procedures for box 2 should provide contact information for each major agency that may control the facility after a disaster.

Box 3 leads to five parallel activities. The first, "Notify Insurance Company," is related to the "Manage Legal and Other Concerns" activity in the beginning sequence. In this case, the various insurance carriers will need to send their own adjusters and engineers to help assess the damage.

The procedures for box 4 should provide the disaster recovery team with specific contact information for all insurance carriers with which the company has policies in force. The insurance company must be given the opportunity to investigate the premises before any salvage, repair, or cleanup operations begin. Unlike the period immediately following a disaster, when some cleanup may be required for safety and health reasons, the insurance company will look very unfavorably at any attempt to disturb the scene prior to it inspections. This does not mean that facility repairs must await a leisurely investigation and inspection by the insurance company; there are immediate steps to take to make preparations for salvage and repair. The procedure for box 4 will aid the disaster recovery team in deciding what they can do immediately, how they can document status prior to repairs, and what must await the approval of the insurance company and others.

Boxes 5 through 7 are the control procedures for salvaging hardware, records, and the facility. The procedures for "Salvage Hardware" must contain all the contact information for hardware vendors, including computer, communications, reproduction, and others. Each type of equipment requires different salvage operations; the best procedure for a computer may, for example, be the worst procedure for a magnetic tape drive. The salvage procedures should provide general guidelines and contact information for quick access to salvage experts. The same is true for facility salvage procedures.

Salvaging vital records is a complex and unpredictable process. For some records, drying is the top priority; for others, getting them into oxygen-free containers is far more important. The procedures supporting the "Salvage Vital Records" activity must provide contact points, general instructions, and guidance for making decisions about salvage possibilities. Unlike hardware, where the salvage decision is usually driven by cost considerations alone, salvaging vital records and legacy materials may be driven by other factors, such as corporate legacy and history considerations or legal issues. The salvage decision is often much more subjective in this area.

Box 8 is labeled "Reconstruction or Relocation." This task differs from the other tasks in three important ways. First, it is likely that the members of the disaster recovery team will not have the required authority to make this decision but will play an important

role in presenting analyses to the decision-making top management people. Second, there are no specific company-external contacts to be made by the disaster recovery team as a part of this process. Although a real estate agent may be contacted for information, it is not likely that the disaster recovery team will undertake that task. The third difference is that the result of this activity is not some specific observable process but a set of recommendations made to top management. The procedures built to support this activity should provide guidance to the disaster recovery team so that they can develop the analyses and present the results in a timely fashion.

Box 9 represents the last phase of the disaster recovery process, initiating the process that returns operations to a normal operating environment. At the time the plan is written, there is no way to predict whether by following this step the company will be relocating to a new facility, or making minor repairs or none, and returning to the existing facility. The importance of these last steps is that they confirm the operability of the newly restored and operational systems. Each of these activities needs to be supported by a set of procedures that guide qualified personnel in their tasks of restoring and verifying systems, networks, data, and records.

59.5 IMPLEMENTATION AND READINESS. The final steps in preparing the disaster recovery plan are proving that the plan works and ensuring that the people understand their roles. This is performed through rehearsals, walk-throughs, and testing. Once the plan is completed and all of the reserve systems and capabilities have been installed, it is time to test whether these backup systems work as planned. There are a variety of testing phases, starting with extremely limited tests performed separately from ongoing operations and proceeding through full tests using live data.

The first set of tests is intended to prove that specific functions and capabilities are in place and available for use. These are the onsite and offsite limited tests, whose goal is to segregate tasks into small groups of related functions. With only a few functions per tests, there can be many tests that, in the aggregate, prove that each required function is successfully represented in the suite of recovery capabilities. The onsite portion of these tests will *disconnect* the users from normal operating capabilities such as LAN, WAN, and support staff.

This is best accomplished by setting up a testing room within the facility, but away from operations. Users are given the tools and instructions they would receive during a real disaster, and they attempt to perform their designated disaster recovery functions. An example of such a function is emergency payroll operations, where the outsource company is notified that this is only a test or the backup company facility is put on notice that a test is in progress. The tests would determine whether offline systems were able to process and transfer data as expected. It also may demonstrate that manual procedures work as written. In this testing, the keys are to set test objectives in advance, monitor test progress, ensure that users have only the expected disaster recovery tools at their disposal and no others, and document the results. After the onsite tests prove satisfactory, these same basic, limited procedures should be followed off site as they would be during an actual emergency. This requires some employees to stay away from the office, so the tests should be planned to avoid periods of expected high activity.

Once the functions have been tested, the next set of tests is intended to prove that the technical systems can be made operational in the planned time frame and that the systems can support users as expected. Two types of tests accomplish this: offsite network tests and hot site tests. These are discussed together because they are often overlapping capabilities performed at the same time. The offsite network and large-

system hot site tests demonstrate that the systems and networks can each be restored within the time allotted and that the interface between them survives the transition to hot site operations. One of the primary advantages of performing the offsite network and hot site tests concurrently is that connectivity and intersystem functionality can be demonstrated.

The last test is a full operations test that attempts to show that all of these elements will work together. Just as in a real emergency, the test will focus on a subset of normal operations. Critical functions are tested first, and other functions are added according to the recovery timeline. For these operations, a test disaster is declared. Employees are notified that this is a test and that they are to cease normal operations and immediately implement their recovery procedures.

Most disaster recovery tests use simulated data or backup data that are real but are not the operational data. Changes made and transactions recorded during such tests are never brought back into the operational environment. Processing performed during the test is based on simulated business operations, with *representative* transactions rather than real ones. In the past decade, there has been a trend toward making final full operations tests online, using actual backup sites, systems, and procedures for processing real transactions. This makes the test process significantly more exciting, with a corresponding increase in stress, for employees and management. Mistakes made in this type of testing can cost the company dearly, but success can assure the company of a recovery plan that is much more likely to succeed during a real crisis.

Employees participating in full operations tests, using real data, are going to be much more serious, but employee rewards are also greater. There is more of a feeling of accomplishment if the test processed live data rather than simulated data. After the test, employees can return to work without needing to catch up on missed work. Management is also more likely to be directly involved in a live test than in a simulated test—because in a real test, management will be called on to make real decisions. They are more likely to be present, just as they would be during a real emergency. Not all organizations are structured in such a way as to make live tests possible. When they can be used, they are the most valid test of the recovery plan. However, they should never be attempted until the plan has been proven using multiple instances of the full complement of tests that should precede the live test.

59.6 CONCLUDING REMARKS. Over the past three decades, the very nature of the threats we considered, and the preparations we made, have changed dramatically. In the early days, major fears included equipment failures and minor geographic events, such as storms and small floods. In the event of a major catastrophe, we felt confident that our plans would unfold as well as those of our neighboring companies; perhaps there would even be some cooperation between companies.

In the decade of the 1990s, however, the criticality of systems combined with an upsurge in man-made disasters forced us to rethink our practices, scenarios, and procedures. The first bombing of the New York World Trade Center in 1993, the bombing of the Murrah Federal Building in Oklahoma City in 1995, the damage to the Pentagon and the heinous destruction of the World Trade Center twin towers in 2001, as well as attacks and failed attempts worldwide since then have finally led us to the point where we must include *artificial* disasters in our planning as well as natural ones. The disaster recovery planner must now consider disasters and recovery scenarios that include terrorism as well as tornadoes. The task is difficult, but the techniques and technologies are the same as before.

59.7 FURTHER READING. See Chapter 58 in this *Handbook* for a list of suggested readings. Additional suggestions follow:

Websites

Disaster Recovery Institute International. *Education*. https://drii.org/education/education.php

Disaster Recovery Journal. www.drj.com

Disaster Recovery News. www.itworld.com/disaster-recovery

Disaster Recovery Overview. U.S. Economic Development Administration. www.eda.gov/disasterrecovery.htm

Disaster Recovery Resources. Clair B. Rubin. www.disasterrecoveryresources.net/aboutus.php

National Disaster Recovery Framework. Federal Emergency Management Agency (FEMA). www.fema.gov/national-disaster-recovery-framework

Books:

Gregory, P., and P. J. Rothstein. *IT Disaster Recovery Planning for Dummies*. Wiley, 2007.

Gustin, J. F. *Disaster and Recovery Planning: A Guide for Facility Managers*, 6th ed. Fairmont Press, 2013.

Hiatt, C. J. *A Primer for Disaster Recovery Planning in an IT Environment*. Information Science Reference, 1999.

López-Carresi, A., M. Fordham, B. Wisner, I. Kelman, and J. C. Gaillard, editors. *Disaster Management: International Lessons in Risk Reduction, Response and Recovery*. Routledge, 2013.

Rothstein, P. J., ed. *Disaster Recovery Testing: Exercising Your Contingency Plan*. Rothstein Associates, 2007.

Schreider, T. *Encyclopedia of Disaster Recovery, Security & Risk Management*. Crucible Publishing Works, 1998.

Stein, M. *Disaster Recovery 81 Success Secrets—81 Most Asked Questions On Disaster Recovery—What You Need To Know*. Emereo Publishing, 2013.

Toigo, J. *Disaster Recovery Planning: Getting to Business-Savvy Business Continuity*, 4th ed. Prentice-Hall, 2013.

Wallace, M., and L. Webber. *The Disaster Recovery Handbook: A Step-by-Step Plan to Ensure Business Continuity and Protect Vital Operations, Facilities, and Assets*, 2nd ed. AMACOM, 2010.

CHAPTER 60

INSURANCE RELIEF

Robert A. Parisi, Jr., John F. Mullen, and Kevin Apollo

60.1 INTRODUCTION 60·2
- 60.1.1 Historical Background 60·2
- 60.1.2 Growing Recognition of the Need for Insurance 60·3
- 60.1.3 General Liability Issues 60·3

60.2 INTELLECTUAL PROPERTY COVERAGE 60·4
- 60.2.1 Recent Developments in Comprehensive General Liability for Intellectual Property 60·4
- 60.2.2 Loss/Damage to Intangible Assets 60·7
- 60.2.3 Intellectual Property Policies 60·8
- 60.2.4 Claims Made versus Occurrence Coverages 60·8
- 60.2.5 Duty to Defend versus Indemnity 60·9
- 60.2.6 Who Is Insured? 60·10
- 60.2.7 Definitions of Covered Claims 60·10
- 60.2.8 Prior Acts Coverage 60·11
- 60.2.9 Extensions of Coverage 60·11
- 60.2.10 Common Exclusions 60·11
- 60.2.11 First-Party Coverage and Other Key Provisions 60·11

60.3 PROPERTY COVERAGE 60·12
- 60.3.1 Direct Physical Loss 60·12
- 60.3.2 Loss of Income 60·12
- 60.3.3 Damage from Employees 60·14
- 60.3.4 Necessary and Reasonable Costs for Recovery 60·14

60.4 CRIME/FIDELITY COVERAGE 60·15

60.5 E-COMMERCE POLICIES 60·15
- 60.5.1 Damage to Third Parties 60·16
- 60.5.2 Direct Financial Loss 60·17

60.6 PRIVACY AND IDENTITY THEFT EXPOSURES 60·17
- 60.6.1 Issues for Businesses 60·17
- 60.6.2 Civil/Legal Liability 60·18
- 60.6.3 Regulatory Issues 60·18
- 60.6.4 Recent Laws Governing Liability for Data Privacy Breaches 60·19

60.7 DAMAGES TO BRAND IMAGE, CONSUMER CONFIDENCE 60·23
- 60.7.1 Special Considerations 60·24

60.7.2	Business Insurance for Privacy Risks	60·24		60.7.5 Insurance for Consumers	60·26
60.7.3	Other Commercial Insurance Policies	60·25	**60.8**	**CONCLUDING REMARKS**	**60·27**
60.7.4	Issues for Consumers	60·26	**60.9**	**FURTHER READING**	**60·27**
			60.10	**NOTES**	**60·27**

60.1 INTRODUCTION. This chapter presents an overview of traditional insurance products, and how they may or may not provide coverage for the risks associated with intellectual property and with computer and network security. It also addresses the new types of coverage that have been developed expressly to address those risks.

60.1.1 Historical Background. Historically, people responded to the risks associated with commerce by finding ways to lessen the impact or severity of potential losses.

- Around 3000 BC, Chinese merchants cooperated by distributing cargo among several ships prior to navigating dangerous waterways, so that one ship's demise would not cause a total loss to any individual.
- Around 1700 BC, Babylonian merchant caravans were constantly imperiled by bandits. In addition to any risk-management practices they might have employed similar to those of Chinese merchants, the Babylonians opted for a risk transfer approach. ". . . Babylonian king Hammurabi developed a code of law, known as the *Code of Hammurabi*, which codified many specific rules governing the practices of early risk-sharing activities. For instance, the code dictated that traders had to repay merchants who financed trading voyages unless thieves stole goods in transit, in which case debts would be cancelled."[1]
- In addition to the well-known Colossus, the people of Rhodes created the *general average*—an insurance construct that survives to this day. The people of Rhodes were a superstitious lot, and would often attribute the fact that a voyage was not going well to some aspect of the cargo being cursed; once identified, the "cursed" cargo was heaved overboard—much to the chagrin of its owner. In response to this practice, which apparently became a burden on trade, Rhodes developed the general average, whereby all the noncursed stakeholders on the voyage contributed *pro rata* to making the unlucky soul whole—at least economically.
- Similarly, as Venice rose to power in the thirteenth century AD and began to venture farther afield in pursuit of commercial and political gain, it encountered such diverse problems as spoilage, pillage, and piracy. The clever Venetians established the practice of pooling their funds, and would then indemnify any losses from that pooled fund—essentially creating the first mutual insurance company.
- Finally, during the late seventeenth century, as international commerce grew, there developed the modern approach to risk transfer, whereby disinterested third parties put up financial capital against the likelihood of a loss of a ship or cargo. The third parties, who frequented Edward Lloyd's coffeehouse, were called *underwriters* because they would literally write their names on posted pieces of paper under the names of the ships they desired to insure.

60.1.2 Growing Recognition of the Need for Insurance. The theft or unintended disclosure of personally identifiable, sensitive, or confidential information led to increased concerns and ever expanding data breach regulations (state, federal, and non-U.S.) and litigation across the United States. First- and third-party liability exposure for data breaches, and related privacy risks, are mounting for all businesses that collect and maintain personal information about customers and/or employees. Individuals and business entities are facing new threats (in addition to ever-present threats of employee and vendor error) that may compromise data stored on computer systems, networks, and paper.

Among the many companies encountering higher-risk attendant to data breach are healthcare organizations, financial services entities, retail establishments, hospitality businesses, and professional service providers. Insurers are encountering data breach risks and cybersecurity concerns both in the context of underwriting and claims handling, given the significant, and ever-increasing, role of data and electronic commerce in the marketplace. The concomitant potential for significant, and costly, breaches of electronic and paper information cause a wide spectrum of public and private entities to recognize the need not only to manage these risks, but to insure against potential losses from them.

Even a perfect network-security infrastructure would be subject to human error. Every organization must integrate the more traditional ways that businesses manage the risks of simply doing business. A company does not decline to buy fire insurance because it installed a new sprinkler system; a bank does not refrain from purchasing crime coverage because it has the best available vaults and guards.

Companies traditionally bought insurance to protect against damage or loss of their principal assets—their plant, their fleet of vehicles, and so on. Where do they turn now that their principal assets are the data and information stored on their business networks?

60.1.3 General Liability Issues. Businesses traditionally managed risks through the purchase of insurance, principally *commercial general liability* (CGL), and property policies. Additional risks faced by financial institutions and professional services firms usually were addressed through a mix of fidelity/surety bonds and specialty insurance products, for example, errors-and-omissions insurance. What they have in common, however, is that all these lines of coverage were created before the widespread, commercial use of computers and networks, especially the Internet and the many statutes addressing data security. Not surprisingly, these traditional policies often fall short of addressing the increasingly complex data and network security exposures that companies face today.

The CGL policy is the most fundamental component of a corporation's insurance portfolio and, not surprisingly, the coverage that receives initial attention in the wake of a data loss or computer breach. Originally called *comprehensive general liability* (also CGL), such policies provide a wide range of coverage for the liability of an insured, because of bodily injury or property damage suffered by a third party as a result of the insured's action or inaction. This coverage, which is explored later in greater detail, is the perceived situs of coverage for intellectual property infringement claims—an issue of critical importance today, when so much of a company's worth, even its existence, is based on its rights under patents, copyrights, or trademarks.

The basic CGL policy provides coverage for all sums the insured is legally obligated to pay as a result of bodily injury, property damage, or personal and advertising injury.

However, the potential bodily injury associated with information technology is quite limited. There is also confusion due to the use of a common term, *privacy*. The CGL policy responds to a claim of *invasion* of privacy, but typically will not respond to a privacy *breach*, absent facts and circumstances that otherwise bring the event into the ambit of the CGL policy's terms and conditions.

60.2 INTELLECTUAL PROPERTY COVERAGE. Of critical importance to a company dealing with new technologies is the potential that its new innovation is a variation on existing and protected intellectual property (as discussed in Chapter 11 in this *Handbook*). In such cases, the innovator may be forced to cease production, to obtain a license from the entity that owns the intellectual property at issue, or to seek to have the legal protections of the applicable area of intellectual property invalidated through the courts.

Alternatively, the company in possession of protected technology, created through the expenditure of millions of dollars for research and development (R&D), may suddenly find that such technology has been incorporated in the new product of a competitor. The victimized company is faced with loss of market share, and possibly with a threat to its very existence. It must enforce its legal property rights—either through a license to the offender, or by seeking to preclude the use of the technology. In both situations, all parties will be incurring significant legal fees to resolve the issue.

In CGL policies, coverage does not always exist (on a defensive basis) for the infringement of a third-party's intellectual property. Such coverage is neither a subject of a broad grant nor a particularly effective one. To trigger coverage, generally, the infringement must be deemed an offense under the policy's *Personal and Advertising Injury* coverage section. Even then, coverage is limited by applicable definitions, exclusions, and the restrictive insuring agreement.

60.2.1 Recent Developments in Comprehensive General Liability for Intellectual Property. Two recent iterations of the *Insurance Services Office, Inc.* (ISO) language for the ISO CGL form CG 00 01 (July 1998) and (October 2001) highlight the significant differences to coverage that may be afforded depending on the policy language at issue. For example, the July 1998 version provides:

> **COVERAGE B. PERSONAL AND ADVERTISING INJURY LIABILITY**
>
> **1.** Insuring Agreement
>
> **a.** We will pay those sums that the insured becomes legally obligated to pay as damages because of "personal and advertising injury" to which this insurance applies. We will have the right and duty to defend the insured against any "suit" seeking those damages. . . .

By virtue of the use of the defined term—*advertisement*—the form is clear that advertising injury liability coverage is applicable solely to offenses committed in the course of advertising the insured's business (i.e., goods, products, or services).

This grant of coverage is conditioned by caveats that are restrictive in the old economy of brick-and-mortar commerce and are all the more troublesome when viewed in context of the Internet and potential Internet- and data-related liability. To understand the coverage that may or may not exist, it is first necessary to understand how the liability policy defines the operative terms.[2]

The July 1998 CGL policy includes the following terms:

> "Personal and advertising injury" means injury . . . arising out of one or more of the following offenses:
>
> f. The use of another's advertising idea in your "advertisement"; or
>
> g. Infringing upon another's copyright, trade dress, or slogan in your "advertisement."

As discussed, the grant of coverage is limited. There is no coverage for an infringement, for example, if it does not satisfy certain standards, such as the three-pronged test utilized in *New Hampshire Ins. Co. v. Foxfire*[3]:

1. Advertising activity of the policyholder;
2. A claim that falls within one or more of the enumerated advertising injury offenses; and
3. Causal nexus between the offense and the advertising activity.

This representative approach is merely one of several threshold requirements that must be satisfied before coverage will apply. The policy form requires that the advertising injury arise out of an enumerated offense and that the claim seeks damages recoverable under the policy. The question of coverage is further limited by several exclusions that serve to drive home the limited scope of coverage.

Of course, the terms of the specific policy at issue play a significant role in evaluating any claim scenario. The 2001 ISO CGL revision both clarifies and restricts coverage with respect to liability for an insured's Internet activities. For example, included in the definition of "advertisement" are notices that "include material placed on the Internet or on similar electronic means of communication;" and "[r]egarding websites, [are] only that part of a website that is about your goods, products, or services for purposes of attracting customers or supporters." In addition, the 2001 ISO CGL form modified and expanded the meaning of "coverage territory" to include personal and advertising injury offenses "that take place through the Internet or similar electronic means of communication."

However, the growth of Internet activity prompted a significant restriction in potential coverage, as the ISO CGL form explicitly states that "electronic data is not tangible property." Under this form, electronic data means "information, facts, or programs stored as or on, created or used on, or transmitted to or from computer software, including systems and applications software, hard or floppy disks, CD-ROMs, tapes, drives, cells, data processing devices, or any other media which are used with electronically controlled equipment." These provisions appear in subsequent versions of ISO's CGL forms.

The courts' disparate construction of the CGL policy form policies created some confusion in the marketplace. The situation worsens when applied to the Internet, depending on the policy language involved. More specifically, an exclusion regarding ". . . an offense committed by an insured whose business is advertising, broadcasting, publishing, or telecasting" makes the coverage particularly problematic.

If the policyholder specifically operates in the media business, such an exclusion's applicability is clear and unequivocal. However, what happens in the more likely scenario when the policyholder is not a multimedia service provider, but merely an e-tailer or an entity that maintains an informational Website? Not surprisingly, there are two schools of thought.

One group, championed by the plaintiff's coverage bar, argues that, unless the policyholder is actively engaged in the business of advertising, broadcasting, publishing, or telecasting as those industries have been traditionally defined, the exclusion is not applicable. This logic rests on the argument that computers and the Internet do not alter the fundamentals of how the world does business.

The second school of thought looks to the way that computers and the Internet blur the line between traditional media activities and a company's online presence. By maintaining any presence online, the policyholder's activities now fall within the parameters defined by the exclusion. This position bears further exploration.

It is rare for a Website not to contain a banner advertisement, hyperlink to a third-party's site, or material imported from sources other than the policyholder's cache of copyrights and trademarks. The presence of such content may place the policyholder in the shoes of an advertiser, broadcaster, or publisher. It is challenging to differentiate between the newspaper or television network that accepts advertisements, and the e-tailer that carries a banner ad.

The applicability of such an exclusion becomes less clear when a business Website carries no banner ads or hyperlinks, and contains no material beyond that which has sprung from the corporate mind of the policyholder. The question then becomes one that does not involve the misuse of a third-party's content. Does the simple maintenance of a Website presume a business is advertising, broadcasting, publishing, or telecasting? The logic behind that position argues that the insured, by maintaining a Website, is effectively sidestepping the services of publishers and broadcasters and doing the job itself.

Finally, another argument considers the manner, speed, and magnitude with which business is conducted via the Internet. The medium of the Internet was not contemplated when the insurance industry drafted some of the initial versions of the CGL policy. Similarly, premiums charged for most CGL policies do not reflect any actuarial data related to the Internet. In fact, the very nature of the Internet lies in direct opposition to most assumptions an insurance underwriter formulates when calculating a premium.

For instance, when a business contemplates launching an advertising campaign, a concept traditionally goes through some preliminary review or in-house analysis. That concept is then taken to an advertising agency, where, among other things, it would be reviewed for general propriety (whether it infringed another's copyright, trademark, etc.). The final product then would be marketed to either the print or broadcasting media, where the publisher would review the advertisement. In granting advertising injury/personal injury coverage, the insurance underwriter assumes the content has gone through similar independent reviews. On the Web, however, such levels of scrutiny are typically nonexistent.

It is unlikely that these arguments will be resolved definitively until more cases reach the courts. Even then, there will be issues surrounding the applicability of boilerplate policy language. Several carriers have either clarified the intent of the policy by expressly excluding some aspect of Internet activity or have simplified it by not offering such coverage to those entities that have an Internet exposure. Several companies have addressed the issue by creating a policy that affirmatively covers the risks associated with a company presenting information on the Internet, be it substantive content or just a banner advertisement.

It should be noted that the 2001 ISO form included several important exclusions to the personal and advertising liability coverage provided. An important exclusion bars coverage for personal and advertising injury arising out of the alleged infringement of copyright, patent, trademark, trade secret, or other intellectual property rights. This

exclusion includes an exception for claims alleging infringement, in an insured's advertisement, of a copyright, trade dress, or slogan. The 2001 ISO form also sets forth an exclusion applicable to Internet communications, precluding coverage for certain personal and advertising injury claims against an insured who maintains an Internet business or an electronic chatroom or bulletin board. In addition to manuscript endorsements, liability policies may include versions of exclusions promulgated by ISO, such as ISO's 2007 standard form CG 00 01 12 07, that apply to (1) personal and advertising injury arising out of the unauthorized use of another's name or product in the course of the insured's Internet activities, or (2) claims resulting from a violation of applicable laws restricting/prohibiting communications involving facsimile machines, telephones, or computers.

60.2.2 Loss/Damage to Intangible Assets. Property damage is generally defined as physical injury to tangible property including all resulting loss of use of that property and loss of use of tangible property that is not physically injured. However, as discussed in Section 60.2, the 2001 ISO CGL form (and subsequent forms) states that electronic data is not tangible property. In addition, the 2004 version of ISO's CGL form restricts coverage in connection with a loss of electronic data, excluding from coverage "[d]amages arising out of the loss of, loss of use of, damage to, corruption of, inability to access, or inability to manipulate electronic data." These 2001 and 2004 revisions likewise appear in subsequent ISO-promulgated CGL forms, such as CG 00 01 12 07.

Several cases have looked at the nexus between the virtual world and property damage. In one case, a large assembler of personal computers purchased disk drives from a manufacturer and subsequently alleged that the disk drives were defective. It did not, however, allege that the defective drives caused any harm to the other components, nor did it claim any loss of use. The manufacturer presented the claim to its CGL carrier, which later denied the claim. In the resultant coverage litigation, the court ruled the claim did not involve property damage because "physical incorporation of a defective product into another does not constitute property damage unless there is a physical harm to the whole."[4] The implication is clear: The mere fact that a piece of hardware or application is defective will not trigger covered property damage. The importance of this distinction cannot be underestimated. This issue will be emphasized further in Section 70.0, which addresses first-party coverage as a result of direct loss.

However, some argue that coverage may exist under certain policies for loss of electronic data, even in the absence of loss of data due to physical injury to tangible property. For instance, ISO promulgated form CG 00 65 12 04 and form CG 00 65 12 07, which afford coverage on a claims-made basis for sums that an insurer becomes liable to pay as damages because of the loss of electronic data, caused by an electronic data incident. The 2007 version of this ISO form sets forth the following definitions:

> "Electronic data" means information, facts, or programs stored as or on, created, or used on, or transmitted to or from computer software (including systems and applications software), hard or floppy discs, CD-ROMs, tapes, drives, cells, data processing devices, or any other media which are used with electronically controlled equipment.
>
> "Electronic data incident" means an accident, or a negligent act, error or omission, or a series of causally related accidents, negligent acts, or errors or omissions, which results in "loss of electronic data."
>
> "Loss of electronic data" means, damage to, loss of, loss of use of, corruption of, inability to access, or inability to properly manipulate, "electronic data."

In addition, in connection with the 2004 revisions to ISO's CGL form, ISO included a CGL coverage endorsement relating to electronic data liability, form CG 04 37 12 04, which affords more limited coverage than the referenced claims-made coverage forms.

Of course, these types of coverages are subject to various restrictions and exclusions, including retroactive dates.

In most instances, the more critical and interesting issue remains whether damage to or loss of computer data is property damage not only in the context of third-party liability coverage, but also first-party property coverage (discussed in Section 60.3). Generally, courts that have looked at the issue have held that loss of data in isolation does *not* constitute damage to tangible property.[5] This conclusion is not universal and is under attack.[6] However, the failure of an insurer to define "tangible property" clearly in a liability policy may raise the risk of a court's finding coverage.[7]

60.2.3 Intellectual Property Policies. As a result of the emerging importance of technology and the increased recognition of the value of intangible assets, various insurance products have been developed to fill the void.

These policies fall within three general types:

1. Third-party liability;
2. Prosecution or abatement; and
3. First-party liability/loss.

The most common forms of insurance coverage available for intellectual property are third-party liability policies, including policies offering errors and omissions or professional liability coverage. In these policies, coverage may exist for an insured's liability to a third party for infringing on that party's intellectual property rights. There is a limitation on this coverage, though. Coverage usually is conditioned on the infringement's being part and parcel of the insured rendering the service that is the basis (or trigger) of the coverage. Additionally, coverage usually is restricted to claims for copyright or trademark infringement, but generally carries express exclusions relating to trade secrets and patents.

This sort of coverage is written either on a claims-made or occurrence basis. Traditionally, firms that provide content (e.g., advertising agencies, media firms, and publishing companies) have found coverage for copyright and/or trademark infringement in multimedia liability policies that are written on an occurrence basis. The policy language generally covers an insured for claims arising out of wrongful acts that occur during the policy period. In contrast, most other errors and omission/professional liability policies are written on a claims-made basis. They provide coverage for claims made and reported to the carrier during the policy period. The two types of policies can be distinguished further by the fact that claims-made policies may provide coverage for wrongful acts of the insured that occurred prior to the policy period, often dating back several years.

Firms looking for more explicit coverage for intellectual property infringement claims, including coverage for trade secrets and patents, usually are left to look to policies with the sole purpose of providing such coverage.

60.2.4 Claims Made versus Occurrence Coverages. Insurance policies generally fall into one of two types: claims-made or occurrence based. The differences in such policies are relatively simple to grasp.

An occurrence insurance policy provides coverage for claims that are made at any time so long as bodily injury or property damage, caused by wrongful, liability-producing acts of the insured, occurred during the policy period.

It generally does not matter if the claim is made long after the policy has expired. Traditionally, some CGL, and most media liability policies, have been written on an occurrence basis; this group of policies is often the first line of defense for claims of copyright and trademark infringement.

A claims-made insurance policy provides coverage for claims made during the policy period, regardless of when the wrongful acts giving rise to the claim occurred, unless the policy imposes a retroactive date. A claims-made policy would respond to a claim made against the policyholder only during the policy period, even if the wrongful acts complained of occurred several years prior to the inception of the policy. Nearly all professional liability/errors and omissions policies, and some hybrid CGL policies, are written on a claims-made basis. Nearly all policies that offer express and specialized coverage for intellectual property infringement liability are written on a claims-made basis.

How the occurrence/claims-made distinctions apply to these data and technology driven exposures is generally of interest only to insurance brokers and underwriters. Some insurance companies initially took the approach that the Internet was entirely a media risk and, as such, should be addressed by occurrence-based coverage. In contrast, other carriers viewed the exposure as entirely based in the services a Website provides, even if those services are media related, and offer only claims-made coverage to Internet businesses. Still another, smaller group of insurance carriers took the approach that the Internet presents risks that are best served by both types of coverage and have offered blended policies.

The benefits of occurrence versus claims-made policies really depends on whether the claimant is likely to discover and assert its claim soon after the wrongful act, and whether the policyholder has any existing insurance coverage that will respond to potential claims from its prior actions.

60.2.5 Duty to Defend versus Indemnity. Those policies providing express coverage for intellectual property are written on both a duty-to-defend and an indemnity basis. The differences between the two types of coverage are less subtle than the claims-made/occurrence distinction, although the two share some of the same history.

Generally, most liability policies written by domestic insurance companies today obligate the insurance company to defend the policyholder so long as the asserted claim alleges facts that might reasonably be expected to result in coverage under the policy. Different jurisdictions take different views on how to interpret if, and when, that duty is trigged and when the duty is no longer owed. The question is the subject of no small amount of case law and legal analysis. For purposes of this discussion, it need merely be noted that the general duty to defend under the law is broader than the insurer's duty to indemnify the policyholder for damages.

Some older policies, as well as those offered through Lloyds' of London, are written on an indemnity basis. The practical effect of this is that the policyholder must incur the cost of defending itself and paying damages for which it is held liable. The insured then must seek indemnity or reimbursement, subject to any retention or deductible, from the carrier.

Whether a policyholder is better served by purchasing an indemnity or a duty-to-defend policy depends not only on its ability to fund a defense but also on its desire to

maintain complete control over that defense. Pursuant to a duty-to-defend policy, the carrier, subject to the deductible or retention, provides a defense for covered claims from day 1, including designating and appointing counsel.

Duty-to-defend policies usually are associated with a pay-on-behalf-of component. This element of coverage obligates the carrier to pay on behalf of the policyholder any covered damages for which the policyholder is held liable, subject only to the applicable deductible or retention.

In the case of an infringement liability policy, the duty to defend is triggered when a claim alleges that an insured has violated or infringed on the intellectual property rights of another in the course of its business or, if more narrowly underwritten, in the course of the expressly designated covered activity. It is important to note that the validity of the claim is often irrelevant to the duty of the carrier to provide a defense.

60.2.6 Who Is Insured? The typical policy generally provides coverage to the principal or named insured, and to subsidiaries of the named insured; the named insured usually is the entity that applied for the insurance and completed the application. Subsidiary coverage traditionally has been limited to actions taken while the entity at issue was actually a subsidiary of the named insured. Coverage is also customarily provided to any present or former partner, officer, director, or employee of the named insured or subsidiary, but only while acting in his or her capacity as such. Often it includes the estates, heirs, legal representatives, or assigns of deceased persons who were insureds at the time that the intellectual property infringement was alleged to have been committed, and the legal representatives or assigns of insureds in the event of the insured's incompetence, insolvency, or bankruptcy.

Persons that may qualify as insureds under a policy are potentially large in number and diverse. Traditionally, intellectual property infringement liability policies have focused not so much on the *who* but on the *what* of coverage. In this case, the *what* is the scope of the policyholder's business activity for which the insurance carrier has agreed to provide coverage. This can involve anything from a single product or service to the entire breadth of the policyholder's operations.

60.2.7 Definitions of Covered Claims. The value of intellectual property coverage can best be viewed by how it defines a claim. As illustrated by current cases, more often than not, the first salvo fired by a plaintiff is not one seeking damages, but rather one seeking to enjoin the infringer from further infringement. Many policies, however, require that the claimant seek monetary damages before coverage attaches. Thus, a company seeking broader coverage should make sure that the policy would respond to actions seeking:

- A demand for money, services, nonmonetary, or injunctive relief; or
- A suit(s), including a civil, criminal, or arbitration proceeding, for monetary or nonmonetary relief.

Integral to a covered claim is how the policy defines claims expenses. Simply offering coverage for those expenses incurred in defense of an infringement action often affords only half of an effective defense. Often the best defense really is a strong offense. In order for defensive coverage to be effective, it needs to incorporate coverage for expenses incurred in seeking to challenge the validity of the patent that a company's product or service is allegedly infringing.

60.2.8 Prior Acts Coverage. Obtaining coverage for prior acts can be crucial when the coverage is written on a claims-made basis. A company must consider how long it has been doing what it is seeking to protect. Ideally, coverage should go back to the inception of the company; unfortunately, such broad coverage is not always available. Carriers often limit coverage to a discrete number of years prior to policy inception—such as for the time when the policyholder instituted certain internal intellectual property controls and/or when it began the activities for which coverage is sought.

Prior acts coverage is not an issue for policies written on an occurrence basis, since, by definition, they provide coverage only for bodily injury or property damage that occurs within the policy period.

60.2.9 Extensions of Coverage. Most intellectual property insurance policies currently available are written to cover only a specific type of intellectual property infringement, chief among them a patent infringement policy. The typical company, assuming that such an entity even exists, does not have the luxury of being sued only for an isolated activity. It is more probable that a company, by its very existence, will be subject to potential exposures across the entire spectrum of intellectual property. As such, companies need coverage that will respond to an allegation of more than just patent infringement.

Several insurance carriers offer coverage options that address this broad range of intellectual property (IP) risks in a single policy form. Other carriers have addressed this issue by endorsing only that coverage that the company specifically requests. The most common coverage extensions available are for copyright and/or trademark infringement.

60.2.10 Common Exclusions. The most common exclusions considered when evaluating third-party liability coverage generally relate to intentional or criminal activities. Several policies contain absolute exclusions that bar coverage when the infringement is willful or intentional. Other carriers soften the impact by providing a defense until the prohibited conduct is proven. Even then it is possible that coverage will apply to those insureds under the policy that did not know of or participate in the willful conduct.

Along similar lines is an exclusion for punitive or exemplary damages. Historically, such damages have been excluded as a matter of course. Recently, however, some carriers have offered coverage for awards of such damages to the extent that to do so is not against public policy, or otherwise against the law.

Other common exclusions track with the principle that liability insurance is meant to cover fortuitous risks. They include claims arising out of breach of contract, antitrust activities, and infringements that existed prior to the inception of the policy. Liability policies are meant to cover unforeseen risks, not known claims or the cost of doing business.

60.2.11 First-Party Coverage and Other Key Provisions. One breed of coverage differs by its very nature. This *first-party coverage* reimburses the policyholder for the loss of the value of its intellectual property after it is declared invalid (patents) or misappropriated (trade secrets). In addition, coverage is available in the market for a policyholder's loss of a trade secret resulting from a computer attack.

At the moment, the number of carriers offering such coverage is very limited, and the underwriting is, in a word, intense. The first coverage mentioned is called *patent*

validity coverage, which indemnifies buyers and/or sellers of patent rights for loss related to the patents subsequently being declared invalid or held unenforceable. Such coverage can be tailored to the premium that the company wants to pay. The basic coverage pays up to the purchase price of the patent rights. Expanded coverage pays for the loss of expected royalty income. This coverage is limited to patents only.

Another recent innovation provides first-party coverage for loss of a company's trade secrets. At least two Internet-focused policies also provide first-party coverage to a policyholder for the loss of its trade secrets. These policies provide coverage for the assets a policyholder has decided to treat as a trade secret. The policies generally require that the trade secrets be misappropriated through some deficiency in the security of the policyholder's computer system.

60.3 PROPERTY COVERAGE. Of particular importance to any commercial entity that values its information is the protection, or lack thereof, afforded under existing property policy forms.

60.3.1 Direct Physical Loss. Traditionally, such policies provide coverage for the direct financial loss suffered by the policyholder for *direct physical loss of or damage to* an entity's physical, covered property as a result of such brick-and-mortar perils, or covered causes of loss, as fire, windstorm, theft, and the like. ISO's "Causes of Loss—Basic Form" CP 10 10 04 02 provides a representative list of covered causes of loss. An *all risk* property policy provides coverage broader in scope, but it is still usually subject to the brick-and-mortar restraint.

The realities of how things work today, and the shift away from a company's worth being largely comprised of physical assets, suggests that such tangible property–based coverage is no longer adequate or sufficient, in and of itself. In today's economy, the lifeblood of an organization is not its buildings and equipment, although the events of September 11 and recent storms demonstrate how vital coverage for such assets is; in fact, a fairly standard business model is for a company to have elements of its infrastructure provided to it only virtually by any number of third-party service providers. The latter often replace many of the systems that a company would have leased or purchased outright only a few short years ago.

Such relationships create further problems when applied to the traditional property policy. The typical policy provides coverage for tangible, covered property at the physical locations of a business. As such, even if one overcomes the hurdle of whether intangible assets are or are not covered property under the policy, it is likely that such property is not resident at the locations covered under the policy.

60.3.2 Loss of Income. Often purchased in conjunction with the standard property policy is coverage for loss of income due to a business interruption. This coverage customarily has the same direct physical loss of or *damage to* trigger, but whereas the basic property policy seeks to indemnify the policyholder for the value of the actual lost or damaged property, business interruption coverage reimburses a policyholder for loss in the form of (1) the loss of net income plus normal operating expenses that continue during the covered period, or extended period, of interruption; and (2) the necessary extra expense incurred during the period of restoration in the effort to continue normal business operations. In addition, coverage often is extended to include an element of reimbursement for contingent or dependent business interruption. This provides coverage for a policyholder if a business that the policyholder depends on is interrupted by a peril covered under the policy, and subsequently causes the

policyholder to suspend business as a result. Some recent property policies, such as policies based on ISO's "Business Income (and Extra Expense) Coverage Form," CP 00 30 04 02, set forth explicit limitations on business income and extra expense coverage when the suspension of operations is caused by an interruption of computer operations (such as from the destruction or corruption of data). In addition, forms such as ISO's "Business Income Form Dependent Properties—Broad Form" endorsement, CP 15 08 06 07, preclude dependent property business income coverage "when the only loss to 'dependent property' is loss or damage to electronic data," and impose other related restrictions on business income coverage as related to such lost or damaged electronic data and other property.

In today's world, most if not all, businesses rely on third parties to maintain some element of their network or computer systems, from the obvious application service provider (ASP), or hosting company (ISP), to the less obvious, but equally vital, outsourcing of network security. The concept of the weakest link applies not only to those allowed within a firewall—such as supply chain elements and extranet members, but also to backbone and infrastructure suppliers. A network or site will be just as inoperative whether the distributed denial-of-service attacks the business directly, or brings down its hosting company or ASP, backup and business continuity plans notwithstanding.

A case in Arizona called into question much of what the insurance industry felt was well settled on the issue of tangible versus intangible property. In *American Guarantee & Liab. Ins. Co. v. Ingram Micro, Inc.*,[8] a policyholder sought business interruption coverage under a property policy for loss of access, use and functionality of its computer systems, which resulted from the loss of custom programming information in the random access memory of three mainframe computers that occurred due to a brief power outage. In finding for the policyholder, this federal court based its decision largely on federal computer crime law and its belief that the continued distinction between *tangible* and *intangible* property, embodied in the policy term "physical damage," was not justified.

Not surprisingly, the *Ingram Micro* case has had a galvanizing effect on the legal and insurance community. The policyholder's bar has hailed it as a commonsense decision; whereas some carriers and some in the legal community at large have pointed at holes in the logic, personal opinion, and leaps of faith in the decision in criticizing the opinion. What all can agree on, however, is that it has created an increasing awareness and sensitivity by companies as to their insurance portfolios. In addition, subsequent case law has thrown the *Ingram Micro* decision into question, rejecting an expansive interpretation of what constitutes physical loss or damage to property. *See Ward Gen. Ins. Servs., Inc. v. Employers Fire Ins. Co.,* 7 Cal. Rptr. 3d 844 (Cal. Ct. App. 2003) (loss of computer data not direct physical loss). In addition, carriers have endeavored to include provisions in policies designed to exclude coverage for computer-related loss. For example, ISO-based exclusions, such as IL 09 35 07 02, exclude coverage for loss or damage due to the failure, malfunction, or inadequacy of computer hardware, software, operating systems, and networks, and any products, services, data, or functions that use or rely on any of those items, due to the inability to correctly process dates or times.

In addition to the efforts of commercial property insurers to manuscript endorsements that exclude coverage for loss of data, ISO introduced a revised "Building and Personal Property Coverage Form," in 2002, CP 00 10 04 02, which provided that electronic data was not covered property, except as provided in an additional coverage (subject to various policy conditions). The most recent version of the ISO "Building

and Personal Property Coverage Form," CP 00 10 06 07, contains provisions similar to the 2002 version of the form. Covered property does not include "electronic data," defined as follows:

> Electronic data, except as provided under the Additional Coverage, Electronic Data. Electronic data means information, facts, or computer programs stored as or on, created or used on, or transmitted to or from computer software (including systems and applications software), on hard or floppy discs, CD-ROMs, tapes, drives, cells, data processing devices, or any other repositories of computer software that are used with electronically controlled equipment. The term *computer programs*, referred to in the foregoing description of electronic data, means a set of related electronic instructions that direct the operations and functions of a computer or device connected to it, which enable the computer or device to receive, process, store, retrieve, or send data.

60.3.3 Damage from Employees. This 2007 ISO form provides additional coverage for the cost to replace or restore destroyed or corrupted electronic data, but the coverage is subject to a $2,500 annual aggregate sublimit, as well as other limitations. In addition, this extended coverage is somewhat limited by its very terms. For example, the 2007 ISO form states as follows:

> The Covered Causes of Loss include a virus, harmful code or similar instruction introduced into or enacted on a computer system (including electronic data) or a network to which it is connected, designed to damage or destroy any part of the system or disrupt its normal operation. But there is no coverage for loss or damage caused by or resulting from manipulation of a computer system (including electronic data) by any employee, including a temporary or leased employee, or by an entity retained by you or for you to inspect, design, install, modify, maintain, repair, or replace that system.

A logical extension of increased interaction and dependencies between service providers and clients, as well as between co-venturers and simple conversants, is that, increasingly, businesses will come to hold, or at least touch, more and more information assets besides their own in their computer system. Such property, like its electronic brethren, tortured judicial logic notwithstanding, will not rise to the level of *covered property*, which generally is limited to tangible property located at, near, in, or on the policyholder's insured premises.

It has been suggested that coverage might be found under a property policy pursuant to a *valuable papers and records* extension. The problem with this argument is that, absent express endorsement to the contrary, valuable papers and records are usually defined as manuscripts and the like, not electronic data or the media used to store or record such data. However, different policy wording may yield different results. *NMS Servs. Inc. v. Hartford,* 62 F. App'x 511, 514-15 (4th Cir. 2003) (hacker erased vital computer files and databases; destroyed property constituted a "record" within the policy definition of "valuable papers and records").

60.3.4 Necessary and Reasonable Costs for Recovery. There also have been attempts to resurrect the notion that coverage might be found under the *sue and labor* aspect of first-party property coverage. This would result in affording coverage not for the actual property loss itself but for the necessary and reasonable costs associated with preventing or fixing the perceived problem. Much was made of this in numerous cases seeking coverage for year 2000 (Y2K) remediation efforts; these cases died quiet deaths due in large part to the *known* element of Y2K (the technological equivalent of a burning building in that programmers were well aware of the problem)

and the nonevent that Y2K eventually proved to be. An analogy in relation to network and computer security would be for an insured to look to its insurance company to pay for the company's firewall and intrusion detection system, or for evaluating the effectiveness of those systems.

It is worth noting that a few carriers have expressly added coverage for loss or damage to intangible assets and resulting business interruption from nonphysical perils. This approach is decidedly in the minority but does highlight that when the markets want to address an issue, they do so directly and expressly.

60.4 CRIME/FIDELITY COVERAGE. Traditional crime and fraud policies have provided, and do provide, a certain level of coverage for direct financial loss due to computer fraud. These policies, however, have usually been limited to indemnifying a policyholder for loss of *money, securities, and other property.* Such policies also can contain limitations as to coverage involving both the intent and the identity of the thief, that is, whether an employee of the policyholder or a third party. However, these policies commonly fall short in offering protection or indemnification for the theft or misappropriation of information, and of intangible assets such as trade secrets, data, and technology—the essential elements of e-commerce.

Those policies that do afford express coverage for the loss suffered by a policyholder as a result of a computer crime generally limit that coverage to the loss of intangible property or information assets. In some instances, courts may focus on rules of insurance contract interpretation that result in a finding of coverage. For example, in *Retail Ventures, Inc. v. National Union Fire Ins. Co. of Pittsburgh, PA*, 691 F. 3d 821 (6th Cir. 2012), the Court of Appeals for the Sixth Circuit held, through a questionable analysis, that a commercial crime policy provided coverage for certain losses resulting from a hacking incident that occurred at DSW in which unauthorized access to a main computer system resulted in the downloading of credit card and checking account information pertaining to more than 1.4 million DSW customers. Construing the policy phrase "resulting directly from" very liberally in favor of DSW, the Court of Appeals ruled that the policy did not "unambiguously limit coverage" to loss resulting solely or immediately from the theft at issue, thereby affording coverage for direct financial loss sustained due to a hacker's infiltration of a main computer system. Time will tell whether this ruling will be generally applied. Not surprisingly, carriers have responded to the decision by adding language to their crime/fidelity wordings that attempt to plug the issues identified by the Court.

Common fidelity policies also can include loss that is other than the deprivation of the information asset, such as situations where the crime involves the copying of data. In today's economy, a company can effectively suffer a loss where its information assets are only copied—for example, the disclosure of a client list, business plan, vendor relationship details, and the like.

60.5 E-COMMERCE POLICIES. In the time since the 5th edition of this *Handbook* was published in 2009, the marketplace for ecommerce policies has matured, with considerably more uniformity among the competing insurance products. One of the things that the markets do agree on is a focus on all industries as opposed to simply concentrating on the technology, telecommunications, and financial sectors. The markets offering policies with definitive coverage for information technology, e-commerce, and network security have begun to grow to meet increasing demand in a technology-driven, increasingly regulated economy, though new entrants now come more from established carriers as opposed to small start-up entities—with over 30 markets now offering some

sort of stand-alone e-commerce or "cyber" policy. What remains important to note is that all of these forms must, and will, continue to change with time as the technology, business practices, and regulations to which they are tied change and evolve. One need only look back to the early 1990s to find policies that spoke of the infinite channels of the Internet, and information security as part and parcel of the general content-related risks faced by publishers and broadcasters.

These e-commerce policies generally fall into two broad categories: (1) those that cover damage to third parties by the actions or failures of the policyholder's computer systems and/or its mishandling of confidential information, and (2) those that cover direct financial loss of the policyholder from certain specified cyber perils. Several carriers also offer loss prevention and loss mitigation services tied to covered claims or losses.

60.5.1 Damage to Third Parties. The most common e-commerce policy is that which would defend and indemnify a policyholder for claims made against it by others for damages allegedly suffered by those third parties. In addition to filling the gaps in traditional policies, such as a CGL policy, such policies also extend new coverage that is peculiar to the needs and exposures of information technology and computer system security.

Typical third-party exposures include the infringement of another's intellectual property, violation of privacy rights, content-based liability such as libel and slander, and professional malpractice. Claims alleging such wrongs are more often than not attributable to an error or act of negligence by an insured. They do not differ dramatically from their old economy cousins, with the exception of the myriad of privacy issues that recently have come to the fore in the form of general privacy, healthcare, and financial information.

The new exposures or risks that these policies respond to include the damage caused by the transmission of malicious code, the unwitting participation in a distributed denial of service attack as a zombie,[9] and others. These perils also can be the triggers for coverage for the direct loss suffered by the policyholder itself in the form of stolen information assets, corrupted data, network interruption, cyber extortion, and so on.

The more robust e-commerce policies have adapted and will continue to adapt to the inherent differences between the Old and the New Economy. Policyholders risk finding themselves bare at worst, or at best succeeding in forcing a settlement after protracted litigation, when they seek to shoehorn coverage for today's risks into hoary old policies from the days of black-and-white television and the Cold War. New coverages have always developed around new exposures, such as employment practices; so too have new coverages arisen to address the new exposures of: a company relying on a cloud service provider for elements of its infrastructure, divergent regulations either requiring proactive action to adopt a standard or imposing a penalty, and/or remedial response in the wake of a defined incident.

As these new e-commerce policies cover new risks, the due diligence performed by the insurance carrier in underwriting the risk must break new ground. The carriers that have taken the time to educate themselves about the risks inherent in computer and network security have turned to computer security professionals to assist in their due diligence. This due diligence generally takes the form of data security assessments, remote penetration testing, a holisitic security assessment, and a full-blown on-site security audit, or some combination. What is important to note is that these activities basically focus exclusively on the applicant's computer and network security issues

and policies; as such, they represent a true value added to the applicant in the form of an additional set of eyes taking a critical look at its security.

60.5.2 Direct Financial Loss. In addition, companies have developed policies and marketed coverage that would allow individuals to cover the damages and costs associated with the loss and repair of their financial or credit identity. Such coverage also has been available to individuals under traditional personal lines policies, such as homeowner's insurance, although that capacity seems to be disappearing as the risks associated with such loss become more widespread and publicized.

Recent events continue to highlight the vulnerability of even the most robust computer, data, and network security. One need only check the news on any given day to learn of a cyber attack from organized crime or hostile state, a loss or mishandling of data by a major corporation, or other event highlighting the truly different way that the world works today. Apart from simply shutting down a business, the only way for a business to operate with the confidence that it will survive a computer attack and to instill that same confidence in its trading partners, be they customers or suppliers, is to have that traditional old risk-transfer vehicle, the insurance policy, sitting behind that business.

The issue of computer, data, and network security is itself changing. Accounting changes, shifts in asset composition, regulation, and the leveraging of intangible assets have caused network and computer security to evolve from a technology issue into a management liability issue. That evolution requires that companies address both the efficacy of their technology and internal policies and the completeness of their insurance portfolio. Of critical importance today is the issue of privacy.

In light of changing economic and regulatory realities, insurers have developed specialty insurance coverage products (using many different naming conventions) that afford consumers and businesses options to address various cyber data risks, including protecting against potential loss resulting from a security data breach. These so-called cyber security insurance policies sometimes offer a wide range of coverage to businesses, for example, offering coverage for third-party claims asserted by consumers affected by a data breach, or for costs incurred in responding to regulatory actions. Carriers are increasingly offering identity theft coverage to individuals, as well as other types of related coverage, as discussed more fully in Section 60.6.3. Cyber data security policies may also afford first-party coverage to protect against damage due to a theft or loss of data assets, cyber extortion, cyber terrorism, or even lost laptops. Of course, these types of policies may be subject to various limitations, restrictions, and exclusions on coverage, as is the case with traditional insurance policies.

60.6 PRIVACY AND IDENTITY THEFT EXPOSURES. Businesses of all sizes across all industries hold both confidential commercial and personally identifiable, information in their care. Both businesses and consumers can take advantage of insurance to mitigate the consequences of an unauthorized disclosure or mishandling of confidential information, including the increasingly troublesome issues created by new and often confusing regulation from all levels of government and self-regulating industries.

60.6.1 Issues for Businesses. Consumers, employees, clients, vendors, merger partners, and other constituents expect the businesses they entrust with confidential information to safeguard it. Importantly, so do legislators and regulators. When a business fails to do so, the consequences can be severe.

60.6.2 Civil/Legal Liability. Companies may be held legally liable for losses that consumers or other third parties incur if information in their care is accessed or compromised, or if identity theft occurs. Even if a business has sound security practices in place, and is not negligent in handling private data, the cost of defending against the allegations and the settlements that often result can be significant. In addition, insurers providing noncyber policies will challenge claims for defense and indemnity coverage in response to a claim arising out of a data breach or hacking incident. Following are few of the many recent scenarios that have put companies at risk:

- In June 2005, victims of a massive compromise involving millions of credit card accounts stolen from payment-processor CardSystems Solutions[10] filed a class-action lawsuit demanding damages. Plaintiffs included not only individuals but also retailers affected by the breach.[11] CardSystems were also charged by the Federal Trade Commission (FTC) and agreed to comply with strict standards for security improvements and audits.[12]
- After a hacker stole debit card information in 2006 from TJX Companies, Inc., owners of T.J. Maxx and Marshalls stores across the United States, approximately a dozen banks had to reissue thousands of new cards to customers.[13] The banks sued the retailer for the cost of reissuance and fraud.[14] This case was deemed the largest security breach in history to that date.[15] A settlement agreement was negotiated in 2008 that offered "vouchers, cash benefits, credit monitoring, identity theft insurance, and reimbursements to eligible people affected by the intrusion(s). TJX will also hold a one-time special event reducing prices 15% for one day at T.J. Maxx, Marshalls, T. J. Maxx 'n More, Marshalls MegaStore, HomeGoods, A.J. Wright, Winners, and HomeSense stores on a future date not yet determined."[16] This incident emphasizes the potential, significant liability that may arise out of a data breach. In that regard, the United States Court of Appeals for the First Circuit permitted banks that issued credit and debit cards to customers affected by this security breach to proceed with various state law claims against TJX, and the processing bank for those credit/debit card transactions on TJX's behalf, based on claims that TJX failed to take adequate measures to protect data security.[17]
- In the past several years, carriers have filed declaratory judgment actions denying any CGL-based obligation to defend or indemnify numerous class action lawsuits (and other claims and potential regulatory matters) arising out of incidents in which cyber attacks resulted in the unauthorized access to and the theft of personal identification and financial information of millions.

60.6.3 Regulatory Issues. Various regulatory bodies may pursue administrative actions against entities experiencing a data or privacy breach. Actions may be brought by the Federal Trade Commission (FTC), Office of Civil Rights (OCR), the Securities and Exchange Commission (SEC), the Commodity Futures Trading Commission (CFTC), state insurance authorities, federal banking agencies, and state Attorneys General. Responding to these investigations require a significant amount of managerial time and legal expenses, and may result in costly settlements. In addition, companies may be forced to undertake lengthy, expensive, and ongoing corrective actions as a result of the investigations and settlements.

Illustrating the impact of a regulatory investigation is the widely publicized case of ChoicePoint, in which a major consumer data broker agreed to pay $10 million in civil

penalties and $5 million for consumer redress in January 2006 to settle FTC charges arising from a data security breach. Under the terms of the settlement, the company was also required to implement new procedures to enhance information security and must undergo security audits by a third party every other year for decades. The breach compromised the personal financial records of more than 163,000 consumers and resulted in at least 800 cases of identity theft.[18]

In another example of the steep liability that a company may face as a result of an FTC investigation, in August 2012, Google agreed to pay a record $22.5 million civil penalty to settle charges that it violated an earlier privacy settlement with the FTC. This particular fine arose out of charges that Google had misrepresented to users of an Internet browser that Google would not place "tracking" cookies on a certain consumers' computers, or serve other targeted advertisements to those users. The FTC charged that Google not only circumvented default blocking settings, but made misrepresentations in that regard, thereby violating a prior settlement that barred Google from misrepresenting the extent to which consumers could control the collection of customer information. This is the largest fine that the FTC has ever issued for a violation of an FTC order.[19]

Regulatory requirements, investigations, and actions can affect a wide variety of industries, involving endless factual scenarios, and resulting in a variety of potential costs and penalties. For example, in January 2013, the FTC reached a settlement with the operator of a leading umbilical cord blood bank in response to charges that the company failed to protect the security of customers' personal information. According to the FTC, the company violated its own privacy policy by failing to use reasonable and appropriate procedures for handling customers' personal information. The company had sustained a security breach during which unencrypted backup tapes, a laptop, a hard drive, and a USB drive, all containing the personal data of approximately 300,000 customers, were stolen from an employee's personal vehicle. As part of this particular settlement, each future violation could result in a civil penalty of up to $16,000.[20]

Other security failures have resulted in the FTC filing a complaint against a global hospitality company (and subsidiaries) arising out of alleged data security failures that led to several data breaches. With respect to one of those incidents, the FTC alleged that inadequate security failures led to fraudulent charges appearing on consumers' accounts, millions of dollars in fraud loss, and the export of hundreds of thousands of consumers' payment card accounts to an Internet domain address registered in Russia.[21] Other companies that have faced FTC charges as a result of alleged security breaches include businesses that market products enabling employers to comply with federal immigration laws.[22]

Most states have enacted laws requiring companies to notify consumers if personal, sensitive, or confidential information has been, or may have been, compromised. These laws dictate under what circumstances a company must give notice, and what method of communication must be used.[23] Although requirements vary from state to state, the costs of legal guidance, forensics, and notification often can run into hundreds of thousands of dollars. As of early 2013, only four states have not enacted any such breach notification laws: Alabama, Kentucky, New Mexico, and South Dakota.

60.6.4 Recent Laws Governing Liability for Data Privacy Breaches. Over the past few years, there have been a number of federal and state initiatives that may govern the possible liability that holders of personally identifiable information may incur as a result of a breach.

60.6.4.1 Executive Order—Improving Critical Infrastructure Cybersecurity. On February 12, 2013, President Obama signed an Executive Order[24] designed to strengthen the cybersecurity of critical infrastructure by increasing information sharing regarding cyber threats (between the Federal Government and U.S. private sector entities). This Order defined the term "critical infrastructure" as a system/asset so vital to the United States that its incapacity or destruction would have a debilitating impact on national security, economic security, or public health or safety. The Order requires the Secretary of Commerce to direct the Director of the National Institute of Standards and Technology to lead the development of a baseline framework to reduce the threat of cyber risks to critical infrastructure, including developing "a set of standards, methodologies, procedures, and processes that align policy, business, and technological approaches to address cyber risks." According to the Order, this framework "shall focus on identifying cross-sector security standards and guidelines applicable to critical infrastructure" and "shall incorporate voluntary consensus standards and industry best practices to the fullest extent possible." In addition, this baseline framework "shall include methodologies to identify and mitigate impacts of the Cybersecurity Framework and associated information-security measures or controls on business confidentiality, and to protect individual privacy and civil liberties." As explained in the Order, the Federal Government will create these standards through consultation with owners and operators of critical infrastructure and other stakeholders, and an open public review and comment process. The establishment of these voluntary standards may provide a critical baseline that may lead to future regulations. However, courts or agencies currently evaluating a particular data breach may choose to consult these security standards, when developed, in order to evaluate and to resolve the alleged liability of a holder of personal information for a security breach.

60.6.4.2 SEC Disclosure Guidance. On October 13, 2011, the U.S. Securities and Exchange Commission, Division of Corporation Finance, released a disclosure guidance document, entitled CF Disclosure Guidance: Topic No. 2, which provides the SEC's views regarding disclosure obligations that relate to cybersecurity risks and cyber incidents.[25] Publicly traded businesses must be cognizant of these guidelines when reviewing the adequacy of disclosures that may relate to security risks and incidents. As emphasized in this guidance document, although current disclosure requirements may not explicitly refer to cybersecurity risks and incidents, federal securities laws are designed to elicit the disclosure of timely, comprehensive, and accurate information concerning risks and events that a reasonable investor would consider important to an investment decision. In addition, the disclosure of material information regarding cybersecurity risks and incidents may be required when necessary to make other required disclosure not misleading. Although various operational and financial considerations will impact the adequacy of any disclosure relating to cybersecurity, registrants should disclose such risks "if these issues are among the most significant factors that make an investment in the company speculative or risky." In that regard, registrants should take into account various factors, such as: (1) prior cyber incidents, and the severity and frequency of those incidents; (2) the probability of cyber incidents occurring, and the quantitative and qualitative magnitude of such an incident occurring; and (3) the adequacy of preventative actions taken to reduce cybersecurity risks in the context of the industry in which the registrant operates. Registrants should also address cybersecurity risks and incidents in presenting management's discussion and analysis of the financial conditions of a company's operations "if the costs or other consequences

associated with one or more known incidents or the risk of potential incidents represent a material event, trend, or uncertainty that is reasonably likely to have a material effect on the registrant's results of operations, liquidity, or financial condition or would cause reported financial information not to be necessarily indicative of future operating results or financial condition." Other instances that may require a registrant to disclose such information include "if one or more cyber incidents materially effect a registrant's products, services, relationships with customers or suppliers, or competitive conditions." Thus, the SEC's guidance highlights the need for publicly traded companies to engage in a critical analysis regarding cybersecurity risk and assessment as a component of obligation relating to financial disclosure under federal securities laws. It is unclear if, and how, this SEC guidance affects the obligations of companies to comply with breach notification laws, including the need to inform shareholders concerning breach notifications by way of statements set forth in audited financial statements, or otherwise. However, depending on the registrant's circumstances, the SEC noted that appropriate disclosures may include a "[d]escription of relevant insurance coverage."

60.6.4.3 HIPAA/HITECH. The Department of Health and Human Services issued its 563 page HIPAA/HITECH final rule designed to strengthen the enforcement of rules promulgated under the Health Insurance Portability and Accountability Act (HIPAA).[26] The Health Information Technology for Economic and Clinical Health (HITECH) Act, enacted as part of the American Recovery and Reinvestment Act of 2009 and signed into law on February 17, 2009, was designed to promote the adoption and meaningful use of health information technology, and modified the authority of the Secretary of the Department of Health and Human Services to impose more stringent civil monetary penalties for violations of HIPAA rules. The recent HIPAA/HITECH modifications had several goals: (1) to strengthen privacy and security protections for individuals' health information established under HIPAA; (2) to modify breach notification rules relating to unsecured, protected health information; (3) to modify HIPAA's privacy rules, strengthening privacy protection for genetic information; and (4) to make certain modifications to improve the effectiveness of HIPAA's breach notification and enforcement rules, thereby increasing flexibility for and decreasing burden on regulated entities. It is anticipated that the effect of this final rule will result in a greater number of breaches that require reporting and in an increase in potential exposure to civil penalties and attendant risk. This final rule was published on January 25, 2013, becomes effective on March 26, 2013, and requires covered entities and business associates to come into compliance by September 23, 2013.

60.6.4.4 Gramm-Leach-Bliley Act. Federal law imposes obligations on financial institutions to ensure the security and confidentiality of personal customer information. Financial institutions include businesses significantly engaged in providing financial products and services. In connection with the implementation of the Gramm-Leach-Bliley Act, the FTC issued a final rule entitled "Standards for Safeguarding Customer Information," commonly referred to as the "Safeguards Rule." The Safeguards Rule requires financial institutions subject to the FTC's jurisdiction to develop a written pre-breach security plan composed of standards establishing administrative, technical, and physical safeguards for customers' information.[27] The FTC adopted the Safeguards Rule in order to "ensure the security and confidentiality of customer records and information; protect against any anticipated threats or hazards to the security or integrity of such records; and protect against unauthorized access to or use of such

records or information that could result in substantial harm or inconvenience to any customer." The Safeguards Rule became effective on May 23, 2003.

60.6.4.5 Massachusetts Data Security Regulations. In addition to the federal government, state governments focus on imposing stringent requirements on those individuals and businesses that collect and retain personal information for purposes of engaging in commerce so as to protect against the unwanted and/or unauthorized disclosure of personal information. For example, Massachusetts enacted regulations, effective March 1, 2010, that established standards for persons and companies who collect and retain personal information about a resident of Massachusetts.[28] These regulations require the development and implementation of a written information security program that outlines administrative, technical, and physical safeguards appropriate to protect personal information. In addition, Massachusetts requires each person that owns or licenses personal information about a resident of Massachusetts (and that electronically stores or transmits that information) to include in the required written, comprehensive information-security program, the establishment (and maintenance) of a security system covering computers, including any wireless system.

60.6.4.6 Federal Statues under Consideration in 2013. Federal legislators have introduced a number of recent initiatives that, if passed, could change the landscape regarding liability and costs for data breaches, and have significant impact on possible liability implications for holders of personally identifiable information. At the time of writing (July 2013), the Library of Congress's *Thomas* system and *govtrack.us* provided the following statuses for some of the more significant publicized proposals[29]:

S.1193

Title: DATA SECURITY AND BREACH NOTIFICATION ACT OF 2012.

Full Title: A bill to require certain entities that collect and maintain personal information of individuals to secure such information and to provide notice to such individuals in the case of a breach of security involving such information, and for other purposes.

Sponsor: Sen. Toomey, Patrick [R-Pa].

Introduced: June 20, 2013.

Latest Major Action: This bill was assigned to a congressional committee on June 20, 2013, which will consider it before possibly sending it on to the House or Senate as a whole.

H.R. 1468

Title: SECURE IT.

Full Title: A bill to improve information security, and for other purposes.

Sponsor: Rep. Marsha Blackburn [R-TN7].

Introduced: April 10, 2013.

Latest Major Action: This bill was assigned to a congressional committee on April 10, 2013, which will consider it before possibly sending it on to the House or Senate as a whole.

S.1408

Title: DATA BREACH NOTIFICATION ACT OF 2011.

Full Title: A bill to require Federal agencies, and persons engaged in interstate commerce, in possession of data containing sensitive personally identifiable information, to disclose any breach of such information.

Sponsor: Sen. Dianne Feinstein [D-CA].

Introduced: July 22, 2011.

Latest Major Action: Feb 6, 2012 Placed on Senate Legislative Calendar under General Orders. Calendar No. 310.

S.1511

Title: PERSONAL DATA PRIVACY AND SECURITY ACT OF 2011.

Full Title: A bill to prevent and mitigate identity theft, to ensure privacy, to provide notice of security breaches, and to enhance criminal penalties, law enforcement assistance, and other protections against security breaches, fraudulent access, and misuse of personally identifiable information.

Sponsor: Sen. Patrick Leahy [D-VT].

Introduced: June 7, 2011.

Latest Major Action: Sep 22, 2011. Died (Reported by Committee).

S.3742

Title: DATA SECURITY AND BREACH NOTIFICATION ACT OF 2010.

Full Title: A bill to protect consumers by requiring reasonable security policies and procedures to protect data containing personal information, and to provide for nationwide notice in the event of a security breach.

Sponsor: Sen. Mark Pryor [D-AR].

Introduced: August 5, 2010.

Latest Major Action: August 5, 2010 Senate Commerce, Science and Transportation committee/subcommittee actions.

Status: Still under consideration.

H.R. 6236

Title: DATA BREACH NOTIFICATION ACT OF 2011.

Full Title: A bill to require Federal agencies, and persons engaged in interstate commerce, in possession of data containing sensitive personally identifiable information, to disclose any breach of such information.

Sponsor: Sen. Dianne Feinstein [D-CA]

Introduced: September 28, 2010.

Latest Major Action: September 22, 2011 Reported by committee.

Status: Died.

60.7 DAMAGES TO BRAND IMAGE, CONSUMER CONFIDENCE. When private information in a company's care, custody, or control is accessed, compromised, or stolen, the company can sustain costs, damage to its reputation, and a decline in

customer and employee confidence. How a company communicates the problem to those affected and manages the aftermath of an incident is critical.

If a company responds to a security breach or a potential security breach with rapid, effective, and accurate communications, and affords proactive, practical assistance for victims, it can generally mitigate damage to its public image. Companies that handle private or confidential data must plan, and prepare to mount, an appropriate response to potential security breaches, including considering the costs such a response will entail.

60.7.1 Special Considerations. All businesses that handle personal information are exposed to privacy- and identity theft–related liabilities. Certain businesses hold especially sensitive information, and have a particularly pronounced risk. They include:

- Any company that holds confidential information under a confidentiality agreement.
- Financial services companies that hold information concerning an individual that would be considered *nonpublic personal information* as defined under Title V of the Gramm-Leach-Bliley Act of 1999 (Pub. L. No. 106-102, 113 Stat. 1338), and the FTC's subsequent final Safeguards Rule. This definition encompasses all information on applications to obtain financial services, such as credit card or loan applications, bank or credit card account histories, and the fact that an individual is or was a customer of the financial institution.
- Healthcare companies and their business associates that handle information that could be considered *protected health information* within the meaning of the Health Insurance Portability and Accountability Act of 1996 (Pub. L. No. 104-191, 110 Stat. 1936) (HIPAA), and the HIPAA/HITECH final rule issued by the Department of Health and Human Services in early 2013. HIPAA defines protected health information to encompass all *individually identifiable health information* held or transmitted by a covered entity or its business associate in electronic, paper, or oral form. The new HIPAA/HITECH rules expand business associate obligations to certain subcontractors of business associates.

60.7.2 Business Insurance for Privacy Risks. Privacy insurance coverage enables businesses to manage the significant exposure associated with privacy breaches and identity theft. The insurance addresses the expenses and liability that can result when confidential information in a company's care, custody, and control is accessed, disclosed, compromised, or stolen. The same markets that provide e-commerce policies have responded with specialized coverage for privacy breaches. In evaluating this coverage, it is important to note that while the e-commerce policies generally require a computer malfunction of some sort, the privacy liability coverage responds to the claim regardless of how the privacy breach occurs—be it a computer glitch, Dumpster® diving, a lost laptop, or a rogue employee walking out the door with a thumb drive or paper file.

A sound insurance program should provide coverage for both claims-related expenses as well as out-of-pocket pre-claims expenses that are typically incurred as part of a privacy event. Coverage should also respond the risks and incidents that are caused by the policyholder's vendors. Termed "vicarious liability" or information holder coverage, privacy insurance has expanded to allow for a trigger that does not require any direct error or action by the policyholder; basically treating the breach occasioned by

the vendor as if caused by the policyholder in order to protect the policyholder—with the carrier reserving the right to seek reimbursement from the actual party who caused the breach.

Coverage is available for the following:

- ***Legal liability and defense costs***. This includes judgments, settlements, and defense fees incurred as a result of a claim or civil litigation initiated by consumers, clients, organizations, or other businesses.
- ***Regulatory action expenses***. This includes legal expenses that result from investigations, negotiation of consent orders, and formal adversarial proceedings instituted by government or regulatory agencies. This is a critical privacy coverage, since the regulatory action will often precede the civil suit. It is in the insured's and insurer's best interests to have a favorable resolution of the regulatory action.
- ***Event Response Expenses:*** The costs incurred to hire forensic and legal experts to determine the cause and extent of the event, as well as any legal obligations that may flow from the forensic findings. This also includes the services of a breach coach, typically a law firm that will act as guide and counsel to the policyholder as it manages through the crisis. This is another element of coverage that is typically triggered long before a claim is alleged by a third party.
- ***Notification costs.*** The costs incurred to notify individuals or businesses whose information was compromised or stolen, including costs of mail, e-mail, telephone, or advertising. This is an interesting coverage enhancement, not generally found in other policy forms. It is, however, of growing importance as relates to current and pending state and federal regulations.
- ***Crisis communications management expenses.*** This includes the cost of public relations counsel to help a company shield its reputation during a pivotal period.
- ***Recovery services for victims***. This includes the costs of education, assistance, credit monitoring, and/or identity restoration for customers, employees, or others whose information was stolen or compromised. Providing these services can help a company to retain customers, maintain employee satisfaction, and prevent costly legal actions in the wake of an incident.

A recent and growing trend is for the carrier to provide elements of pre-claim expenses such as event response expenses, recovery services, and notification costs as part of a service offering, included as part of the coverage but sitting apart from the policy's aggregate financial limit as opposed to a simple monetary limit designated for the specific coverage elements. Such an event or breach response service offering approach has proven attractive to small and medium-sized entities, who often lack the sophistication or relationships to respond to an incident. The ability to simply call on the panel of experts assembled by the carrier enables the policyholder to manage the event and mitigate potential loss in a way that they could not do unaided.

60.7.3 Other Commercial Insurance Policies. Many companies operate under the risky and possibly inaccurate perception that traditional insurance policies, such as commercial general liability (CGL), directors & officers (D&O) liability, or errors and omissions (E&O) insurance, will adequately address their cyber, privacy and identity theft exposures. However, traditional commercial insurance policies are generally not designed to address risks related to computer networks and electronic

information. Citing just one example, CGL policies typically do not define coverage *territory* to align with the global nature of the Internet, and do not contemplate the various privacy laws worldwide. These policies would not normally encompass coverage for expenses associated with victim notification, crisis management, and identity theft recovery either.

In fact, company management and other professionals who rely solely on traditional insurance will be surprised to find they have little or no insurance when a privacy breach or identity theft incident occurs. The specialized privacy and identity theft insurance discussed previously was designed expressly to fill this significant coverage gap.

60.7.4 Issues for Consumers. Identity theft, defined as the theft and fraudulent use of an individual's personal identification information, including a Social Security number, account numbers, or other personal data, is the fastest growing financial crime in America. One of every eight adults, or a family member, has been a victim of identity theft.[30] More than 9.3 million individuals were victimized in the United States in 2004.[31]

Identity theft is not only widespread; it is costly. Individuals use unlawfully obtained personal information to purchase goods and services, or to obtain new mortgages, lines of credit, or additional credit cards. A criminal may use a stolen identity to commit employment fraud or to escape criminal prosecution.

According to statistics gathered from consumer complaints filed with numerous state and federal organizations, and reported by the FTC, consumers reported 990,242 fraud-related complaints in 2011. Consumers reported paying over $1.5 billion in those fraud complaints. Categories of identity theft included: government documents/benefits (27%), credit card fraud (14%), utilities fraud (13%), bank fraud (9%), employment fraud (8%), and loan fraud (3%).[32] Medical identity theft is also increasing.

60.7.5 Insurance for Consumers. Consumers in the United States are well protected from the direct financial loss that can result from identity theft. If unauthorized credit card charges and electronic banking transactions are reported promptly, an individual has little financial liability. Still, individuals whose identities are stolen must undertake a time-consuming and often costly process to restore their names and credit. It can take six months to detect identity theft, and up to 600 hours to recover from the crime.[33] In some cases, individuals must pay to dispute fraudulent debts and accounts opened by an identity thief.

Insurance can help individuals alleviate the costs associated with the compromise of their personal information and identity theft. Such insurance is primarily geared to mitigate the costs of recovering from such incidents and to ease the recovery process. A typical policy might provide:

- **Expense Reimbursement,** encompassing numerous fees and expenses required to recover from identity theft
- **Income Protection,** which pays wages lost as a result of time off from work required to recover from identity theft
- **Identity Restoration Services,** including the services of a personal case manager to handle identity recovery work on the victim's behalf
- **Loss Prevention and Mitigation Services,** such as access to a customer service center to assist in preventing identity theft, and credit monitoring to promote early detection of problems

Many financial institutions, associations, employers, and service providers make this type of insurance available to customers, members, and employees.[34] Identity theft insurance may also be available to consumers in conjunction with homeowners' or business owners' policies.

60.8 CONCLUDING REMARKS. The significant exposures resulting from intellectual property and privacy breaches that companies and individuals face today have become serious boardroom concerns. To mitigate against these exposures, a comprehensive insurance program should be a critical component of the risk-management strategies put in place. While insurance for information systems exposures was once considered novel and limited in scope, it has quickly evolved, and new products in the marketplace offer broad and effective protection.

60.9 FURTHER READING

Dionne, G. *Handbook of Insurance*, 2nd ed. Springer, 2014.

Fordney, M. *Insurance Handbook for the Medical Office*, 12th ed. Saunders, 2011.

Ostrager, B. R., and T. R. Newman. *Handbook on Insurance Coverage Disputes*, 16th ed. Aspen Law & Business, 2013.

Sutcliffe, G. S. *E-Commerce Insurance and Risk Management*, 2nd ed. Standard Publishing, 2000. *Out of print*.

60.10 NOTES

1. Norma L. Nielson, "Insurance," *Microsoft® Encarta® 2008* [DVD] (Redmond, Washington: Microsoft Corporation, 2007).
2. Few cases have addressed coverage under common policy formulations defining the term "personal and advertising injury" in the context of an underlying data breach. The Court of Appeals for the Ninth Circuit held that a personal injury liability offense provision, covering claims alleging certain violations of a person's right of privacy, afforded coverage for underlying lawsuits complaining that an Internet provider interrupted and intentionally disseminated private online communications. *Netscape Commc'ns Corp. v. Federal Ins. Co*., 343 F.App'x 271 (9th Cir. 2009). *Compare Novell, Inc. v. Federal Ins. Co*., 141 F. 3d 983 (10th Cir. 1998) (software company claimed that general liability insurer failed to defend a prior action brought by an independent software developer; creation of a similar software program was not an offense committed in the course of advertising goods, products, or services). *Zurich Am. Ins. Co. v. Fieldstone Mortgage Co*., Civil No. CCB-06-2055, 2007 WL 3268460 (D. Md. Oct. 26, 2007), *reconsideration denied*, 2008 WL 941627 (D. Md. Mar. 24, 2008) (underlying complaint alleged that a mortgage company improperly accessed and used customer credit information, without permissible consent, in order to solicit refinance business).
3. 820 F. Supp. 489 (N.D. Cal. 1993); *see also Zurich Ins. Co. v. Sunclipse, Inc*., 85 F. Supp. 2d 842, 852 (N.D. Ill. 2000), *aff'd,* 241 F.3d 605 (7th Cir. 2001).
4. *Seagate Technology, Inc. v. St. Paul Fire & Marine Ins. Co*., 11 F. Supp. 2d 1150, 1157 (N.D. Cal. 1998) (citation omitted).
5. *Rockford Pharmacy, Inc. v. Digital Simlistic, Inc.*, 53 F. 3d 195 (8th Cir. 1995).
6. *See, e.g., NMS Services, Inc. v. Hartford*, 62 F. App'x 511 (4th Cir. 2003) (first party coverage; software development company sought insurance coverage after a hacker erased vital computer files and databases from an internal computer network; the

company sustained "damage to its property" in the form of damage to computers); *State Auto Prop. & Cas. Ins. Co. v. Midwest Computers & More*, 147 F. Supp. 2d 1113 (W.D. Okla. 2001) (third-party coverage; data stored on a computer disk or tape was not "tangible property" within the meaning of a business owners' liability policy; however, the lost use of computers constituted a loss of use of tangible property within the loss of use of property damage coverage afforded); *see also America Online, Inc. v. St. Paul Mercury Ins. Co*., 207 F. Supp. 2d 459 (E.D. Va. 2002) (third-party liability case; no duty to defend Internet service provider under CGL policy for loss of tangible property, as computer data, software, and systems were not tangible property; however, alleged loss of use of computers constituted property damage under a CGL policy that defined property damage as "physical damage to tangible property of others, including all resulting loss of use of that property.") *aff'd* 347 F. 3d 89 (4th Cir. 2003).

7. *Eyeblaster, Inc. v. Federal Ins. Co.,* 613 F. 3d 797, 802 (8th Cir. 2010) (the underlying complaint alleged repeatedly the "loss of use" of a computer; general liability and errors and omissions policies both provided coverage relating to underlying claims that consumer's computer was infected by spyware).
8. No. 99-185 TUC ACM, 2000 WL 726789 (D. Ariz., Apr. 18, 2000).
9. Zombies are computers that have been taken over by an attacker and used to transmit large volumes of traffic. See Chapters 18, 20, 30, and 41 in this *Handbook*.
10. R. Lemos, "MasterCard Warns of Massive Credit-Card Breach," *SecurityFocus*, June 17, 2005, www.securityfocus.com/news/11219
11. J. Evers, "MasterCard Data Breach: Lawsuit Demands Damages," *Silicon.com*, July 7, 2005, http://software.silicon.com/security/0,39024655,39150141,00.htm (URL inactive).
12. Federal Trade Commission, "CardSystems Solutions Settles FTC Charges," FTC Website, February 23, 2006, www.ftc.gov/opa/2006/02/cardsystems_r.shtm
13. T. Lewis, "TJX Sued for Loss of Consumer Data," *ConsumerAffairs.com*, January 29, 2007, www.consumeraffairs.com/news04/2007/01/tjx_folo.html; original article now replaced with M. Huffman, TJX Pays $9.75 Million to Settle Data Breach," *ConsumerAffairs,com*, July 27, 2009, www.consumeraffairs.com/tjmaxx-data-breach
14. R. Lemos, "New England Bankers Sue TJX for Breach," *SecurityFocus*, April 26, 2007, www.securityfocus.com/brief/490
15. J.Evers, "Revealed: World's Largest Security Breach," *Silicon.com*, March 30, 2007, www.silicon.com/retailandleisure/0,3800011842,39166613,00.htm (URL inactive).
16. *In re TJX Cos. Retail Sec. Breach Litig.*, 584 F. Supp. 2d 395, 398 (D. Mass. 2008) (although the court struck from the settlement the provision regarding the one-day special event sale); The TJX Companies, Inc. and Fifth Third Bancorp., Case No. 07-10162.
17. *In re TJX Cos. Retail Sec. Breach Litig*., 564 F. 3d 489 (1st Cir. 2009); *Anderson v. Hannaford Bros. Co.,* 659 F. 3d 151 (1st Cir. 2011) (hackers stole up to 4.2 million credit and debit card numbers of grocery store customers; consolidated multidistrict litigation); *Pisciotta v. Old Nat'l Bancorp*, 499 F. 3d 629 (7th Cir. 2007) (putative class of consumers commenced a lawsuit against a bank based on the alleged failures of the bank to secure adequately personal information; bank's motion for judgment on the pleadings granted).

18. Federal Trade Commission, "ChoicePoint Settles Data Security Breach Charges; to Pay $10 Million in Civil Penalties, $5 Million for Consumer Redress: At Least 800 Cases of Identity Theft Arose From Company's Data Breach," FTC.gov, January 26, 2006, www.ftc.gov/opa/2006/01/choicepoint.shtm
19. Federal Trade Commission, "Google Will Pay $22.5 Million to Settle FTC Charges it Misrepresented Privacy Assurances to Users of Apple's Safari Internet Browser," FTC.gov, August 9, 2012, www.ftc.gov/opa/2012/08/google.shtm
20. Federal Trade Commission, "Cord Blood Bank Settles FTC Charges That It Failed to Protect Consumers' Sensitive Personal Information," FTC.gov, January 28, 2013, http://ftc.gov/opa/2013/01/cbr.shtm
21. Federal Trade Commission, "FTC Files Complaint Against Wyndham Hotels For Failure to Protect Consumers' Personal Information," FTC.gov, June 26, 2012, www.ftc.gov/opa/2012/06/wyndham.shtm
22. Federal Trade Commission, "FTC Settles Charges Against Two Companies That Allegedly Failed to Protect Sensitive Employee Data," FTC.gov, May 3, 2011, www.ftc.gov/opa/2011/05/ceridianlookout.shtm
23. Court decisions provide guidance on such issues. For example, the California Supreme Court held that a customer's ZIP code constitutes "personal identification information" under the Song-Beverly Credit Card Act, prohibiting merchants from requesting a credit card holder's ZIP code. *Pineda v. Williams-Sonoma Stores, Inc*., 246 P. 3d 612 (Cal. 2011); *but see Feder v. Williams-Sonoma Stores, Inc*., Civ. No. 2:11-03070 (WHW), 2011 WL 4499300 (D.N.J. Sept. 26, 2011) (oral request for a ZIP code as part of a credit card transaction did not violate New Jersey's Truth-in Consumer Contract, Warranty and Notice Act).
24. Executive Order: *Improving Critical Infrastructure Cybersecurity*, 2013 WL 503840 (February 12, 2013).
25. See www.sec.gov/divisions/corpfin/guidance/cfguidance-topic2.htm
26. 78 *Fed. Reg.* 5566 (January 25, 2013).
27. 67 *Fed. Reg*. 36484 (May 23, 2002).
28. 201 Mass. Code Regs. 17.00, *et seq.*
29. Library of Congress THOMAS system for tracking legislation: http://thomas.loc.gov
30. K. Gilpin, "Identity Theft Affects Millions, Survey Shows," *New York Times*, September 3, 2003.
31. M. T. Van Dyke, ed., "2005 Identity Fraud Report," Javeline Strategy and Research, January 2005.
32. *Federal Trade Commission, Consumer Sentinel Network Data Book January-December 2011* (February 2012) (Executive Summary).
33. L. Foley and J. Foley, "Identity Theft: The Aftermath 2003," Identity Theft Resource Center, September 2003.
34. M. E. Kabay, "Defending against Identity Theft: LifeLock," *Network World Security Strategies*, March 4, 2008, www.networkworld.com/newsletters/sec/2008/0303sec1.html and "Defending against Identity Theft: Identity Guard," *Network World Security Strategies*, March 6, 2008, www.networkworld.com/newsletters/sec/2008/0303sec2.html

CHAPTER 61

WORKING WITH LAW ENFORCEMENT

David A. Land

61.1 INTRODUCTION 61·1

61.2 GUIDELINES 61·3
- 61.2.1 Relevant Laws 61·3
- 61.2.2 Plan Ahead 61·3
- 61.2.3 Activities You Should Always Report 61·3
- 61.2.4 Activities You Should Not Report 61·6
- 61.2.5 Terrorism 61·6

61.3 LAW ENFORCEMENT RESOURCES 61·7
- 61.3.1 Federal Bureau of Investigation 61·7
- 61.3.2 U.S. Postal Inspection Service 61·8
- 61.3.3 U.S. Secret Service 61·8

61.4 MEMORANDUM OF AGREEMENT 61·8

61.5 HANDLING EVIDENCE AND THE CHAIN OF CUSTODY 61·8
- 61.5.1 Never Work from the Original Media 61·8
- 61.5.2 Issues of Liability 61·9
- 61.5.3 Ask Law Enforcement To Give Back 61·10
- 61.5.4 The Knock at the Door 61·10
- 61.5.5 Keeping Your Operation Running during an Investigation 61·10
- 61.5.6 Using Log Records 61·11

61.6 INFORMATION SHARING 61·12
- 61.6.1 InfraGard 61·12
- 61.6.2 Information-Sharing Initiative—Implementation Plan 61·15
- 61.6.3 Additional Considerations 61·17

61.7 CONCLUSION 61·18

61.8 FURTHER READING 61·18

61.9 NOTES 61·18

61.1 INTRODUCTION. The year 2012 must be recognized as the "Year of Intrusions." Threats to your organization may involve external threats from hackers, crackers, phreakers, and the like, or an internal threat involving a trusted employee(s) or a competitor attempting to steal, acquire, or otherwise damage information critical to your organization's viability. There appear to be two emerging trends. The first is politically driven attacks on organizations, where the attacker views an organization as being complicit in an event contrary to the criminal hackers' (sometimes called *hacktivists*) views. This is especially true where privacy on the Internet is concerned or viewed as being threatened. The other is financial motivation by individual(s) who

have a plethora of agendas to support their illegal activity. Fraud, waste, and abuse by an individual who has legitimate and authorized access to organizational networks or computing systems is also big business for both foreign and domestic entities seeking to gain an intellectual or economic advantage.

Working with law enforcement at all levels, both nationally and internationally, is an important aspect of computer security and your collective need to protect your sites, your sites' information, and your organizations' enterprise infrastructure. The cybercrime paradigm has dramatically shifted to one in which we need law enforcement and they need us. The success of both sides not only depends on but demands collaboration and sharing.

Understanding the needs and constraints of law enforcement before, during, and after the commission of a cybercrime can significantly enhance an organization's opportunity to come back online quickly, with potentially little or no disturbance to users or customers. In contrast, failing to know your counterparts may ensure the attacker's success and escape from successful prosecution. Conveying your needs and the needs of your clients to law enforcement *before* an incident occurs will serve everyone well.

However, working with law enforcement is not an opportunity to assume the role of law enforcement. You must know your organization's limitations and at what point to engage your law enforcement contacts. It is equally beneficial to know the abilities of law enforcement, especially where it applies to your local law enforcement agency. In short, not all local law enforcement agencies are created equal when dealing with cybercrimes, and contacting the proper organization on the first call is critical. Involvement and coordination with your C-level leadership as well as legal and Human Resource staff are all critical players in any decisions made relative to sharing data.

Most organizations, whether they support the corporate or the government sector (e.g., .mil, .gov, .org, .com, or .edu), will at some point have a need to meet, or otherwise collaborate with, a local, state, or federal law enforcement or an intelligence agency. Intelligence agencies of the United States and of other nations are included here to remind the reader that there are 16 U.S. government entities, such as the Department of Homeland Security (DHS), Federal Bureau of Investigation (FBI), and other government departments within the United States, that have an intelligence role tied directly to computer incidents or cyber-related investigations. Furthermore, as cyber-related crimes know no geographical bounds, there is also the possibility that investigative agencies of other nations, such as the Royal Canadian Mounted Police, the German BSI ("Bundesamt für Sicherheit in der Informationstechnologie," also known as the Office for Security in Information Technology),[1] and Interpol, might need to become involved. Where investigations take on an international nexus, the FBI, the Department of Homeland Security, or Department of State will likely be one of the lead organizations conducting the investigation. The need for collaboration with any of these agencies is driven by the inevitable reality that at some point, a computer or cyber-related crime will be committed where the computers or information of an organization are the targets of unwelcome cyberattacks.

The goals of law enforcement, both foreign and domestic, as well as the various U.S. intelligence agencies, tend to vary based on jurisdiction and the mission of the investigating organization. Where a computer or cybercrime may be of little or no prosecutorial value to one agency, another may find there is significant value in pursuing an investigation. Keep that in mind when making your first contact. Likewise, there is absolutely nothing wrong with selling your company's desire to pursue an investigation to law enforcement or intelligence agencies. In some instances, it may be necessary for

you to do some of the initial investigative legwork to make it more attractive for law enforcement to pursue. Much will depend on the nature of the crime and any losses your organization may have experienced. While the FBI states that they are concerned only with crimes that exceed $5,000, the reality is that the threshold for the FBI is much higher, especially given their current workload and the number of investigations they are currently working. Never assume that the FBI is the first, best organization to call. While they maintain investigative purview in many criminal areas, there are a multitude of other organizations, some of which may actually trump the FBI.

To address these inevitabilities, it is prudent to examine your organization for potential vulnerabilities and to develop a logical path forward in addressing any identified weaknesses. In the planning process, identify your supporting law enforcement organizations and meet with them in order to put in place some form of memorandum of agreement (MOA). This memorandum should clearly articulate the responsibilities of each entity; creating this document will save everyone a substantial amount of time, avoid potential miscommunications, and ensure that organizational information and assets are handled in an appropriate manner. As part of this agreement, make it clear what you as the impacted organization expect as part of any investigation, and that reciprocity is expected. Where there is a need to deem information as classified, intelligence organizations can facilitate the clearance of individuals to support the sharing of any evidence found. MOAs are discussed in detail in Section 61.4.

61.2 GUIDELINES

61.2.1 Relevant Laws. In examining computer crime laws, it is worthwhile to understand which laws, and which law enforcement agencies, are responsible for responding to you or your organization. It is equally necessary to know when the need for law enforcement involvement is warranted and unwarranted. When does a computer crime break a federal, state, or local law? Exhibit 61.1 lists crimes that should be reported whenever they are discovered or when there is any uncertainty as to their status as a crime. Exhibit 61.2 presents a list of activities you should not report.

61.2.2 Plan Ahead. When you know when and whom to call, dealing with any type of significant cyberincident can be manageable. Leaving things to chance or uncertainty or assuming that someone else is responsible and will take care of things can leave your organization in a very precarious situation. Know who is responsible for the conduct of a computer-related investigation and, where possible, train for such likelihoods. These efforts can greatly reduce the time from discovery to prosecution. Lastly, keep your senior management as well as your legal team in the loop regarding all contacts with law enforcement as well as what your intended plans are should the need arise.

61.2.3 Activities You Should Always Report. One checklist for activities you should always report is available online at http://i.i.com.com/cnwk.1d/i/tr/downloads/home/computer crime reporting checklist.pdf

61.2.3.1 Report Intrusions or Attacks on Major Networks or Networks that Deal with Sensitive or Classified Data. If sensitive data such as client financial information, medical records, customer credit card information, Social Security numbers, or other personally identifiable information are compromised, you should report it to the authorities. This is also true if the company has government/

Report intrusions or attacks on major networks, or those that deal with sensitive data.

If sensitive data such as client financial information, medical records, customer credit card information, social security numbers, and the like has been compromised, you should report it to the authorities. This is also true if the company has government/defense contracts or deals with other types of regulated information. The FBI's computer crime squad investigates major network intrusions and network integrity violations. You can report these types of attacks to federal, state, and local authorities and let them sort out the jurisdictional issues.

Report intrusions or attacks that result in large monetary losses.

The amount of monetary loss often determines whether a theft is considered a misdemeanor or felony. Felony offenses will get more attention from law enforcement agencies when they involve a bank or most any other type of financial institution. The following site enumerates the laws relative to financial institutions and the requirement to report: http://www.fdic.gov/news/news/financial/1997/fil97124.html

Report cases of suspected state-sponsored or Industrial espionage.

If an intruder goes after your company's trade secrets, or there are indications that individuals within your organization are the targets of a foreign intelligence service, this is a serious federal offense that will be investigated by the FBI who maintains the charter for all domestic counterintelligence investigations, and by other government and military agencies.

Report cases involving child pornography.

This is an offense that is taken very seriously by law enforcement. If child pornography was discovered on your company computers and was not promptly reported, you, as network administrator, may be implicated or held liable in a civil lawsuit. The possession of child pornography images on a computer, no matter who is in possession of the computer, is breaking the law and is subject to prosecution.

Report e-mailed or other digitally transmitted threats.

All states have laws against threatening and harassing communications. Physical threats against individuals, terrorist threats, bomb threats, blackmail, and similar electronic communications should be reported to security personnel and local police.

Internet fraud

If one of your users is a victim of "phishing" scams or other fraudulent activities perpetrated by an e-mail or unknown individuals on the Web, this information can be reported to one of many different Federal Agencies. Within the US you can report to the Federal Trade Commission, the FBI, the US Secret Service, the US Postal Service, the National Fraud Information Center, etc. A complete list of US Government and non-Government agencies as well as investigative organizations of other countries can be found at http://www.consumerfraudreporting.org/governments.php

Report suspected terrorist activities.

If you suspect that your network is being used for communications between terrorists, first and foremost do not attempt to investigate yourself. Report suspected activities to the U.S. Department of Homeland Security, or the FBI via the FBI's "tips" Website.

Exhibit 61.1 Activities That Should Be Reported

defense contracts or deals with other types of regulated or classified information. The FBI's computer crime squad investigates major network intrusions and network integrity violations. The Department of Homeland Security also has a vested interest in both national and international cybercrime, and should be contacted via their toll-free number, 1-866-DHS-2-ICE, or via their Cyber Crimes Website.[2] You can report these types of attacks to federal, state, and local authorities and let them sort out the jurisdictional issues, with the exception of classified information. This is directly reportable to the FBI, as well as the organization's intelligence and security personnel.

Don't report port scanning and similar "nonintrusive" activities.

Although port scanning is often a precursor to intrusion or attack, in most jurisdictions it's not, in itself, a crime. It's more like walking down a hallway in an apartment building and trying each door to see if it's locked. If you find an unlocked door and go inside, that is criminal trespass— but as long as you don't go in, you haven't committed a crime.

Don't report viruses, Trojans, worms, and spyware—at least, not to law enforcement agencies, although there are exceptions.

Malicious software is a huge and ever-growing problem that does a great deal of damage and costs companies millions of dollars. However, law enforcement agencies generally don't respond to individual malware reports. While those who release viruses and other malware can be prosecuted under Title 18 of the U.S. Code, prosecutors generally go after developers of malware that is widely distributed and causes a large amount of harm. If you encounter a new variety of malware, check the pages of popular anti-virus vendors and report to them if it isn't already listed. Remember that the sender of a virus often doesn't even know he or she is sending it. However, if you have evidence that a particular person wrote and intentionally released a piece of malware with the intent to cause harm or otherwise destroy data, you should contact local law enforcement or your local FBI cybercrime squad.

Exhibit 61.2 Activities That Should Not Be Reported

61.2.3.2 Report Intrusions or Attacks That Result in Large Monetary Losses. The amount of monetary loss often determines whether a theft is considered a misdemeanor or felony. Felony offenses will get more attention from law enforcement agencies when they involve a bank or most any other type of financial institution. The following site enumerates the laws relative to financial institutions and the requirement to report: www.fdic.gov/news/news/financial/1997/fil97124.html

61.2.3.3 Report Cases of Suspected State-Sponsored or Industrial Espionage and Potential Terrorism. If an intruder goes after your company's trade secrets, or there are indications that individuals within your organization are the targets of a foreign intelligence service activity, this is potentially a serious federal offense. The FBI or other government and military agencies that may have a vested interest in the matter should be contacted. The FBI maintains the charter for all domestic counterintelligence investigations; however, other government agencies may have an equal stake and thus should be notified. The most important part to take away from this subsection is to limit the number of people who are aware of any incident regarding the potential loss or theft of classified data!

61.2.3.4 Report Cases Involving Child Pornography. Child pornography is a criminal offense taken very seriously by law enforcement at all echelons. If child pornography is discovered on your company computers and is not promptly reported, you, as network administrator, may be implicated or held liable as part of the criminal investigation. The possession of child pornography, which can include images, or other forms of data found on a computer, no matter who is in possession of the computer, is a crime, and is subject to federal prosecution. Once discovered, this data should not be viewed by anyone and immediately reported to a federal law enforcement agency.

61.2.3.5 Report E-Mailed or Other Digitally Transmitted Threats. All states have laws against threatening and harassing communications. Physical threats against individuals, terrorist threats, bomb threats, blackmail, and similar electronic

communications should be reported to security personnel, local police, and the Department of Homeland Security.

61.2.3.6 Internet Fraud. If one of your users is a victim of a "phishing" scam or other fraudulent activities perpetrated by an e-mail or unknown individuals on the Web, this information can be reported to one of many different federal agencies. Within the United States you can report to the Federal Trade Commission, the FBI, the U.S. Secret Service, the U.S. Postal Service, the National Fraud Information Center, and so on. A complete list of U.S. Government and non-Government agencies, as well as investigative organizations of other countries, can be found at www.consumerfraudreporting.org/governments.php

61.2.4 Activities You Should Not Report. http://i.i.com.com/cnwk.1d/i/tr/downloads/home/computer crime reporting checklist.pdf (URL inactive).

61.2.4.1 Do Not Report Port Scanning and Other Similar "Non-Intrusive" Activities. Although port scanning is often a precursor to intrusion or attack, in most jurisdictions it's not, in itself, a crime. It's more like walking down a hallway in an apartment building and trying each door to see if it's locked. If you find an unlocked door and go inside, that is criminal trespass—but as long as you don't go in, you haven't committed a crime. You can, however, collaborate with other similar organizations, and, using many different tools, you can develop a picture of what the scans may mean, and what information is being sought. This is where information sharing across similar organizations such as major banking institutions can be beneficial. It can be done in such a manner as to protect company information while still putting together potential threat vectors.

61.2.4.2 Don't Report Viruses, Trojans, Worms, and Spyware—At Least, Not To Law Enforcement Agencies, Although There Are Exceptions. Malicious software is a huge and ever-growing problem that does a great deal of damage and costs companies millions of dollars annually. However, law enforcement agencies generally don't respond to individual malware reports. While those who release viruses and other malware can be prosecuted under Title 18 of the U.S. Code, prosecutors generally go after developers of malware that is widely distributed, and causes a significant amount of harm. If you encounter a new variety of malware, check the pages of popular anti-virus vendors and report to them if it isn't already listed. Remember that the sender of a virus often doesn't even know he or she is sending it. However, if you have evidence that a particular person wrote and intentionally released a piece of malware with the intent to cause harm or otherwise destroy data, you should contact local law enforcement, the DHS Computer Crimes Unit, or your local FBI cybercrime squad. NOTE: Most FBI cybersquads will be located in major cities, and will not be found in local Resident Agencies, while DHS has an online presence for reporting viruses and malware.

61.2.5 Terrorism. If you suspect that your network is being used for communications between terrorists, first and foremost do not attempt to investigate this incident yourself. Report suspected activities to the U.S. Department of Homeland Security or the FBI via the FBI's "tips" Website, or through the DHS Website.

They can also reduce the potential of having agencies involved that would otherwise not have a need to know (e.g., the local press).[3]

- Local/State Law Enforcement: Call your local police department, county sheriff's agency. Do not call 9-1-1. Ask for the agency's high-tech crimes unit or, its smaller investigation division.
- FBI Computer Crimes Squad: nccs@fbi.gov or 202-324-9164
- FBI Tips site: https://tips.fbi.gov
- U.S. Secret Service Form 401—Cyber Threat/Network Incident Report: www.secretservice.gov/net.intrusionforms.shtml
- Internet Crime Complaint Center: www.ic3.gov
- National White Collar Crime Center (NW3C): www.nw3c.org
- FTC Scams and Identity Theft Website: www.consumer.gov/section/scams-and-identity-theft

61.3 LAW ENFORCEMENT RESOURCES

61.3.1 Federal Bureau of Investigation. The Federal Bureau of Investigation (FBI) is responsible for the investigation and prosecuting of more than 200 federal statutes. Relative to the cyberworld, the FBI has two primary responsibilities.

First, it is the lead law enforcement agency for investigating cyberattacks by foreign or state-sponsored espionage or terrorism. Second, the FBI and DHS also work to prevent criminals, sexual predators, and others intent on malicious destruction from using the Internet and online services to steal from, defraud, and otherwise victimize citizens, businesses, and communities.

The FBI retains jurisdiction over many types of cases, which include: national security, including counterintelligence and counterproliferation; terrorism; and organized crime. The U.S. Secret Service retains jurisdiction where the Treasury Department is victimized or whenever computers are attacked that are not under the purview of the FBI. In certain federal cases, the Customs Department, the Commerce Department, or military organizations such as the Air Force Office of Investigations, Naval Criminal Investigative Service, U.S. Army Intelligence, or U.S. Army Criminal Investigation Division, may also have jurisdiction.

In the United States and its territories, a number of federal laws protect against attacks on computers, misuse of passwords, electronic invasions of privacy, and other transgressions. The Computer Fraud and Abuse Act of 1986 is the statute governing most common computer crimes, although prosecutions may take place under other laws. The Computer Abuse Amendments Act of 1994 expanded the 1986 act to address the transmission of viruses and other harmful code.

Many of the states within the United States have adopted their own computer crime laws.

According to the Federal Investigative Guidelines, federal law enforcement can gather only proprietary information concerning an incident in four ways:

1. Request for voluntary disclosure of information
2. Court order
3. Federal grand jury subpoena
4. Search warrant

61.3.2 U.S. Postal Inspection Service. The U.S. Postal Inspection Service is the federal law enforcement investigative arm of the U.S. Postal Service. Postal inspectors enforce over 200 federal laws in investigations of crimes that may adversely affect or fraudulently use the U.S. mail, the postal system, or postal employees. Many fraud schemes that originate over the Internet (such as auction fraud or multilevel marketing schemes) or that involve payment or delivery via the U.S. mail are under the jurisdiction of the Postal Inspection Service.

61.3.3 U.S. Secret Service. The U.S. Secret Service began as a force investigating counterfeit currency; it still enforces all laws relating to "the counterfeiting of obligations and securities of the United States," which includes financial crimes, identity theft, computer fraud, and any computer-based attacks on the infrastructure of the United States. It is worth noting that the U.S. Secret Service is another investigative agency like the FBI. Both agencies work very well together and collaborate on many cyber-related investigations.

61.4 MEMORANDUM OF AGREEMENT. A memorandum of agreement (MOA) between any organization and any law enforcement agency is a prudent measure. With an MOA in place, both sides of any investigation have a clear understanding of what to expect. Within the MOA, specific points enumerate the breadth and depth of responsibilities.

For those interested, the Department of Justice maintains a list of recent computer crimes cases and those involved in the litigation.[4]

Know your limitations. Know when to call law enforcement and when to deal with a situation yourself. Discuss the inherent problems on both sides, such as public disclosure of a confirmed intrusion and potential liabilities, which may follow any disclosure.

61.5 HANDLING EVIDENCE AND THE CHAIN OF CUSTODY. Many different cyber-related investigative agencies and groups can be found on the Web. For 12 years, the author has been a member of the International Association of Computer Investigative Specialists (http://iacis.org), which offers to the public a comprehensive set of directions relative to computer forensics and the collection of electronic evidence. A step-by-step process for the examination and collection of electronic evidence follows.

61.5.1 Never Work from the Original Media. Prior to any media examination, a bit-for-bit image should be made of the media. Never work from the original media.

1. Examine the media, logically and systematically, by starting where the data of evidentiary value is most likely to be found. These locations will vary depending on the nature and scope of the case. Examples of items to be noted might include:
   - If the media is a hard drive, the number and type of partitions should be noted.
   - If the media is an optical disc, then the number of sessions should be noted.
   - File systems on the media should be noted.
   - A full directory listing should be made to include folder structure, filenames, date/time stamps, logical file sizes, and so on.
   - Installed operating systems should be noted.

- User-created files should be examined using native applications, file viewers, or hex viewers. This includes such files as text documents, spreadsheets, databases, financial data, electronic mail, digital photographs, sound, and other multimedia files, and the like.
- Operating system files and application created files should be examined, if present. This includes, but is not limited to, boot files, registry files, swap files, temporary files, cache files, history files, and log files.
- Installed applications should be noted.
- File hash comparisons may be used to exclude or include files for examination.
- Unused and unallocated space on each volume should be examined for previously deleted data, deleted folders, slack space data, and intentionally placed data. Previously deleted filenames of apparent evidentiary value should be noted. Files may be automatically carved out of the unallocated portion of the unused space based on known file headers.
- Keyword searches may be conducted to identify files or areas of the drive that might contain data of evidentiary value and to narrow the examination scope.
- The system area of the volume (i.e., FAT, MFT, etc.) should be examined and any irregularities or peculiarities noted.
- Examination of areas of the media that are not normally accessible, such as extra tracks or sectors on a floppy disk, or a host-protected area on a hard drive, may be required.
- To facilitate examination of data, including user settings and device and software functionality, the computer may be booted using either a copy of the boot drive or by using a protected program to determine functionality of the hardware and/or software.
- The forensic software used during the examination should be noted by its version, and should be used in accordance with the vendor's licensing agreement. The software should also be properly tested and validated for its forensic use by the examiner or the examiner's agency.

2. At the conclusion of the examination process, provide sufficient notation of any discovered material of an apparent incriminating or exculpatory evidentiary nature.
3. Provide sufficient documentation of all standard procedures and processes initiated, as well as detailed notation of any variations made to the standard procedures.
4. Properly mark any output of the recovered data with appropriate identifiers in accordance with policies from the examiner's agency.

61.5.2 Issues of Liability. In examining any media or conducting any cybercentric investigation, consideration should be given to these areas:

- Proprietary information
- Business-sensitive information
- Export-controlled information
- Downloading copyrighted music or videos
- National defense information

61.5.3 Ask Law Enforcement To Give Back. In many situations where law enforcement is notified of an event requiring further investigation, the agency may ask you for everything even remotely connected to the investigation. At the same time, most law enforcement agencies have the ability to give back. They can provide training to key personnel and information on various topics, including identity theft, espionage, equipment theft, access control, and so on.

61.5.4 The Knock at the Door. What do you do if your first indication of a possible security incident is a law enforcement agent appearing at your office with a search warrant? In many cases, nothing could be more unnerving. However, if you have done your homework and you have an established rapport with this law enforcement agency, you have little to fear. You already have a plan in place for just such a situation, and you know the rights and obligations.

Although, under the best of circumstances, stress can still occur, it is important to remain unemotional and to express complete willingness to comply with the search warrant team. Avoid being defensive or giving the appearance that you are trying to hide anything. Offer all assistance possible and take these seven steps:

1. Read the warrant carefully and understand what the search warrant team wants to search or seize.
2. Notify both upper management and the legal department of current events.
3. If the authorities want to search a local machine, the impact on the organization may be minimal. Get the appropriate manager or supervisor to evaluate the impact and to assist the officers in securing the machine. In cases of child pornography, the police will take the entire computer. Work with them to see if backup copies can be made of critical data and programs prior to removal.
4. If a server is the target, see if copies of relevant portions of the hard drives will satisfy the requirements of the warrant. Where it is not possible to do this, involve your server and network administrators to bring down the server in an orderly fashion. Your plan should also include bringing online any backup server.
5. Rarely will an entire network have to be taken down. Have your network administrator work with law enforcement to ensure orderly access to parts of the network described in the search warrant.
6. After the search and seizure, the team will leave you "an officer's return" regarding what property was seized. Retain this document for review by your legal staff. Make arrangements with the team to follow up with them on the status of your equipment. Offer whatever technical assistance to law enforcement that your legal department deems advisable.
7. Meet with your management, legal counsel, and technical staff after the team leaves to assess the impact of the seizures on your operations. Also, consider what actions need to be taken for continued operations.

61.5.5 Keeping Your Operation Running during an Investigation. Any investigation that involves a continued law enforcement presence requires time and resources of the organization to ensure support to clients or customers. It is also important to anticipate such an occurrence and to develop a plan on how to cooperate with law enforcement without shutting down your operation. In most instances, a computer crime specialist can generate data dumps, make secure copies of files, and create

logs without carting away all of your organization's computers. Administrators should however, have a plan in place to allocate requisite resources in a manner that will avoid, or at least minimize, disruptions in business operations. The following are some suggestions that may help to minimize the disruption caused by an onsite investigation:

Schedule meetings with key organizational members of your information technology organization, as well as other key legal and management employees, and develop a plan that can be implemented should a significant event take place during either normal or noncritical work periods. Once your plan is developed, *test it!*

If any part of the network needs to come down in support of investigative activities, arrange, where feasible, for this outage to take place during nonpeak hours. Your law enforcement counterparts may be of a different mind, with a greater sense of urgency to take down suspect systems. If possible, work out such details beforehand.

Depending on the operating system environment your organization works in, you will need to know where to acquire the necessary logs and audit files. These files should be identified now rather than during the heat of a significant investigation. In taking this extra step early on, you will help minimize the amount of staff time necessary to retrieve the files needed.

Make sure law enforcement does not take software or database files for which no copy or backup exists. Prior to this, and as a matter of policy, ensure that all organizational employees and contractors are following required backup and copying procedures to avoid any unnecessary loss of vital data. Should any law enforcement agency show up with a warrant or subpoena, you are compelled to provide any, and all, information clearly specified within those documents.

Keep the investigation compartmentalized, and on a need-to-know basis. This action will reduce the amount of staff involved and protect the continuity of the investigation. The more people cognizant of an ongoing investigation, the greater the possibility information central to the investigation may be compromised.

Keep the lines of communication open with law enforcement. Knowing about the needs for records or access to information systems in advance can save staff time and system downtime. In short, do not wait for law enforcement to visit you; visit them, and do it frequently. This will establish good rapport and create trust that may be needed at some future time.

Encourage onsite copying of memory and any magnetic media rather than their removal. Computer forensics specialists possess this capability (also known as *fly-away kits*), and not all cases require the removal of equipment to a forensics lab. With law enforcement's cooperation, you can schedule these procedures during nonpeak hours and keep your equipment on the premises and in service.

61.5.6 Using Log Records. Paper records, video surveillance tapes, proximity card records, and any other means of tracking employee activity can all play a significant role during the investigation of computer crimes and incidents. Admittedly, no one wants to assume the role of "Big Brother," but at the same time, an organization's livelihood may well depend on these methods of verifying access. No one metric exists that clearly demonstrates the true nature of insider threats. It is, however, reasonable to believe that greater than half of all computer crimes are committed by current or former organizational insiders. To assess, and perhaps better to detect, this threat to organizational computing systems, internal documentation and monitoring become important evidence trails.

For example, if video cameras monitor the entrances to the computer processing center or labs, the video surveillance tapes document the entrance and exit of any

employee at a given time. If an entrance uses access cards, the reports or databases containing the access records document employee traffic into sensitive areas. The use of conventional access tracking is not by itself foolproof and will not, for example, preclude "piggybacking," which is the act of closely following another individual and gaining entry without proper authentication. Contradictions between access and video records can often document anomalies or possible misdeeds that will require further investigation. For example, if the access control report indicates John Jones entered the computer lab at 6:45 p.m. on 07-12-08, but the videotape for same time shows Sam Smith entering, there is clear indication that Sam has used John's card or access code. The question now remains: Why? Is there a need to call in law enforcement? At first blush, there is no apparent need to make such a call. Only after some initial investigation on the part of your security staff will this perhaps be necessary.

Paper files—personnel records, departmental documents, project logs, programming modification records, sign-out logs for software, and job assignment records—all tell a story. When you want to know who worked on what, or who had access to which project, paper records can often provide the needed history. Audit logs such as those stored on a UNIX or UNIX-type operating system within /var/log can provide even greater illumination. Using the same type of operating system, a system administrator can simply type "last", which will provide login information to that system. With this said, no effort should be made to verify this information on the original media. Doing so will modify files that may be potentially critical evidence. Thus, an original (bit-for-bit) image of any magnetic media should be made, and all investigation conducted on the imaged data and not the original media.

61.6 INFORMATION SHARING. After 9/11, all law enforcement officers will agree that the sharing of information between any organization and law enforcement is critical to the success of cyber-related investigations. Working with any law enforcement agency has both benefits and consequences. Benefits accrue because:

- Most law enforcement agencies, as well as government and commercial entities have cyberinvestigative capabilities. Many of the federal agencies have taken steps to train and educate investigators to ensure that the most comprehensive investigation is conducted with the least amount of disturbance to the victim agency.
- Law enforcement agencies are also being funded to establish, or better train officers in, cyberinvestigative disciplines that, in most instances, facilitate the acquisition of personnel, equipment, and training.
- Working with law enforcement agencies can serve as a conduit to training and education for on-site first responders and cybersecurity personnel. The FBI or any other law enforcement agency benefits from better educating you as to what it wants and needs, as this helps to ensure that all evidence is properly collected and remains intact. Exhibit 61.3 is a draft example of how a memorandum of understanding (MOU) or an MOA between two or more entities might be constructed for the purpose of sharing law enforcement data.

61.6.1 InfraGard. One of the resources in the United States for forming good relations with law enforcement agencies is the InfraGard. The regional meetings of InfraGard members offer an excellent opportunity for sharing information with trusted

EXHIBIT 61.3 Information-Sharing Example

Information-Sharing Initiative—Implementation Plan

Participating Sites

I. Objective

Implement a program to retrieve information of varying degrees of criticality from specified contributors for the purpose of establishing a repository of related data. This is to enable single search capability for the production of a single composite record. The record will be available as a tool to enhance operational needs of the participants. The record should support tactical operations as well as strategic, investigative, and analytical operations.

II. Participating Agencies

This should include, state, local, federal, and commercial law enforcement, investigative, and analytic entities.

Specific and detailed memorandums of understanding (MOU) will be entered into for each cooperating entity. The MOUs will specify the particular information and data for submission, the limitations or restrictions placed on the data by the subscribing entities, and the roles, responsibilities, and facts agreed upon. Each MOU will recognize and acknowledge the oversight and leadership responsibilities exercised by [your company or the agency taking lead responsibility].

III. Security and Location of Equipment

The system's "Data Warehouse" and IS equipment, including but not limited to computers, servers, printers, and peripherals, will be physically located in a secure facility on your location. The building should be a controlled and "Limited Access" facility, populated with data from three levels (described further below) of information from detailed investigative and analytical assessments.

The security of the system will be controlled and accredited in accordance with the "Organization's" Master Security Plan. (Your site here) is a tenant of the Master Plan. Systems maintenance (updates, backups, etc.) and password control will be regulated by the "Organization's" technical security personnel, who will produce quarterly reports articulating the usage and administration of the system.

The "Organization's" Accepted Use and Security banner will be displayed on all Web pages and sign-in screens.

Access to the system will be restricted to member agencies and their described representatives; the total to be determined by the Executive Management Team. Access levels will be defined as follows:

Level 1: Federal, State, and Local Law Enforcement (L/E) Information

This category is primarily concerned with the identification of persons and normally relates to public safety. Included are sensitive but unclassified law enforcement information such as police and related investigative and field reports, arrest and conviction data, identification databases, and the like.

Level 2: Analytical Access Data

Users in this category will have access to all unclassified tactical incident information available from typical federal government sources for the purpose of analyses based on the data that have been manually collated. Examples of this level of data are 1-94 Arr/Dpt Records, Visa Database accesses (Department of State), 1-20 Student and Visitor records, Custom's Treasury Enforcement Communications System (TECS) data, unclassified federal investigative data (Foreign Access Central Tracking System/FACTS and Personnel Access System).

(*Continued*)

EXHIBIT 61.3 (*Continued*)

Information-Sharing Initiative—Implementation Plan

Level 3: Classified and Security-Related Data Sources

Data not generally shared unless the inquiring entity maintains a clearance and a "need to know" the information. These data are normally derived from isolated databases in the control of the federal or contractor organization. Access to these systems will undergo a much higher screening and approval process. Examples of this level of data are the identity-sensitive documentation and classified systems within the FBI.

Sharing of Level 3 data is an additional phase for implementation after the initial system is operational. Submission of data concerning Levels 1 and 2 can be shared, based on the existing capabilities of joined database systems. However, Level 3 data loading presently can only be accomplished by stand-alone disk loading and by manual transport (sneaker net).

IV. Information to Be Shared

Each participating Agency will provide current data relating to Intelligence, Counterintelligence, Counterterrorism, Terrorist Matrixes, Threat Assessments, Law Enforcement (L/E) alerts, and other L/E sensitive activity reports. All participants will be specifically tasked with providing the elements of Homeland Security effecting Energy, Banking and Finance, Agriculture, Food, Water, Public Health, Information Technology and Telecommunications, Transportation and Border Security, Postal and Shipping, Defense, Government and Emergency Services, and Maritime Security.

Queries of the system by participating agencies will enable access of data on violent criminal gangs, information on persons in jail or prison, searches of criminal records, and other data related to criminal activity (arrest, citation, and booking information, and field and investigative reports). The Federal and DOE data available for review will consist of similar and related alerts, State Department Visa application data, Department of Commerce, Homeland Security data and assessments, as well as portions of the Personnel Access System (PAS) and defined Export Control data. Other initiatives include work authority statements and internal security reports concerning suspected attempts to breach security.

V. Proposed Date of Implementation

The time frame of implementation is best approximated to be a minimum of 6 months from date of contract.

VI. Administration Management

The Lead Agency [to be determined] will be responsible for overall management and administration of the system. Each individual agency will be responsible for day-to-day supervision and oversight of the review, as required to ensure the appropriate level of security of data submitted. Likewise, each participating site will ensure proper data integrity as outlined in MOU. This implies that each site will be responsible for ensuring that the data they upload is not infected with any form of virus or Trojan. While this may not be the most appropriate place for insertion, it must be covered.

The Lead Agency will coordinate with the other agencies the standards and compliance protocols, rules of access and operation, policies for dissemination of analytical reports, and security standards, and will specifically conduct periodic and cursory reviews, not less than monthly, to ensure appropriateness of participation and usage of system. Periodic audits will be conducted to ensure that the highest level of quality information is received, and that appropriate retrieval queries are being administered.

(*Continued*)

EXHIBIT 61.3 (*Continued*)

Information-Sharing Initiative—Implementation Plan

A management team of no less than the Lead Agency and (X) member agencies will be formed to execute the fundamental oversight operations described herein. Each participating agency will be represented by its Executive Manager (CIO or Organizational Investigator, Chief of Police, Sheriff, Commander, etc.) for all operations and decisions. Standards and protocols should be specified in any accompanying MOUs. However, full agency members will enjoy several perks. Most important is the notion of one agency, one vote, which covers development and structure processes relative to the oversight principals and directives.

The management team will define reporting processes to evaluate the program on a quarterly basis to ensure that system integrity is intact and to be cognizant of additional improvements that can be undertaken to provide a more thorough and user-friendly product and process.

members of critical infrastructure organizations and with law enforcement officials. The organization describes itself this way:

> InfraGard is an information-sharing and analysis effort serving the interests and combining the knowledge base of a wide range of members. At its most basic level, InfraGard is a partnership between the Federal Bureau of Investigation and the private sector. InfraGard is an association of businesses, academic institutions, state and local law enforcement agencies.[5]

61.6.2 Information-Sharing Initiative—Implementation Plan. The following sections provide guidance through examples of how one can prepare useful MOAs to foster appropriate collaboration between organizations and law enforcement agencies.

61.6.2.1 Objective. Implement a program to retrieve information of varying degrees of criticality from specified contributors for the purpose of establishing a repository of related data. This is to enable single search capability for the production of a single composite record. The record will be available as a tool to enhance operational needs of the participants. The record should support tactical operations as well as strategic, investigative, and analytical operations.

61.6.2.2 Participating Agencies. This should include, state, local, federal, and commercial law enforcement, investigative, and analytic entities. Specific and detailed MOUs will be entered into for each cooperating entity. The MOUs will specify the particular information and data for submission, the limitations or restrictions placed on the data by the subscribing entities, and the roles, responsibilities, and facts agreed upon. Each MOU will recognize and acknowledge the oversight and leadership responsibilities exercised by your company or the agency taking lead responsibility.

61.6.2.3 Security and Location of Equipment. The system's data warehouse and IS equipment, including but not limited to computers, servers, printers, and peripherals, will be physically located in a secure facility on your location. The building should be a controlled, limited-access facility, populated with data from three levels (described further below) of information from detailed investigative and analytical assessments.

The security of the system will be controlled and accredited in accordance with your organization's Master Security Plan. Your site is a tenant of the Master Plan. Systems maintenance (updates, backups, etc.) and password control will be regulated by the organization's technical security personnel, who will produce quarterly reports articulating the usage and administration of the system.

The organization's Accepted Use and Security banner will be displayed on all Web pages and sign-in screens.

Access to the system will be restricted to member agencies and their described representatives; the total to be determined by the Executive Management Team. Access levels will be defined as follows:

61.6.2.3.1 Level 1: Federal, State, and Local Law Enforcement (L/E) Information. This category is primarily concerned with the identification of persons and normally relates to public safety. Included are sensitive but unclassified law enforcement information, such as police and related investigative and field reports, arrest and conviction data, identification databases, and the like.

61.6.2.3.2 Level 2: Analytical Access Data. Users in this category will have access to all unclassified tactical incident information available from typical Federal Government sources for the purpose of analyses based on the data that have been manually collated. Examples of this level of data are I-94 Arr/Dpt Records, Visa Database accesses (Department of State), I-20 Student and Visitor records, Customs' Treasury Enforcement Communications System (TECS) data, and unclassified federal investigative data (Foreign Access Central Tracking System/FACTS and Personnel Access System).

61.6.2.3.3 Level 3: Classified and Security-Related Data Sources. Data not generally shared unless the inquiring entity maintains a clearance and a "need to know" the information. These data are normally derived from isolated databases in the control of the federal or contractor organization. Access to these systems will undergo a much higher screening and approval process. Examples of this level of data are the identity-sensitive documentation and classified systems within the FBI.

Sharing of Level 3 data is an additional phase for implementation after the initial system is operational. Submission of data concerning Levels 1 and 2 can be shared, based on the existing capabilities of joined database systems. However, Level 3 data loading presently can only be accomplished by stand-alone disk loading and by manual transport (sneaker net).

61.6.2.3.4 Information to Be Shared. Each participating agency will provide current data relating to intelligence, counterintelligence, counterterrorism, terrorist matrixes, threat assessments, law enforcement (L/E) alerts, and other L/E sensitive activity reports. All participants will be specifically tasked with providing the elements of Homeland Security affecting energy, banking and finance, agriculture, food, water, public health, information technology and telecommunications, transportation and border security, postal and shipping, defense, government and emergency services, and maritime security.

Queries of the system by participating agencies will enable access of data on violent criminal gangs, information on persons in jail or prison, searches of criminal records, and other data related to criminal activity (arrest, citation, and booking information, and field and investigative reports). The federal and DOE data available for review will

consist of similar and related alerts, State Department Visa application data, Department of Commerce, Homeland Security data and assessments, as well as portions of the Personnel Access System (PAS) and defined Export Control data. Other initiatives include work authority statements and internal security reports concerning suspected attempts to breach security.

61.6.2.3.5 Proposed Date of Implementation. The time frame of implementation is best approximated to be a minimum of 6 months from date of contract.

61.6.2.3.6 Administration Management. The Lead Agency [to be determined] will be responsible for overall management and administration of the system. Each individual agency will be responsible for day-to-day supervision and oversight of the review, as required to ensure the appropriate level of security of data submitted. Likewise, each participating site will ensure proper data integrity as outlined in MOU. This implies that each site will be responsible for ensuring that the data they upload is not infected with any form of virus or Trojan. While this may not be the most appropriate place for insertion, it must be covered.

The Lead Agency will coordinate with the other agencies the standards and compliance protocols, rules of access and operation, policies for dissemination of analytical reports, and security standards, and will specifically conduct periodic and cursory reviews, not less than monthly, to ensure appropriateness of participation and usage of system. Periodic audits will be conducted to ensure that the highest level of quality information is received, and that appropriate retrieval queries are being administered.

A management team of no less than the Lead Agency and (X) member agencies will be formed to execute the fundamental oversight operations described herein. Each participating agency will be represented by its executive manager (CIO or organizational investigator, Chief of Police, Sheriff, Commander, etc.) for all operations and decisions. Standards and protocols should be specified in any accompanying MOUs. However, full agency members will enjoy several perks. Most important is the notion of one agency, one vote, which covers development and structure processes relative to the oversight principals and directives.

The management team will define reporting processes to evaluate the program on a quarterly basis to ensure that system integrity is intact and to be cognizant of additional improvements that can be undertaken to provide a more thorough and user-friendly product and process.

61.6.3 Additional Considerations. There is, however, another side to any relationship with law enforcement that must be considered. Despite your efforts to be as forthcoming as possible with law enforcement, situations may arise that will cause you and your organization great angst. There may be times when an investigator is simply not able to fully disclose all information gleaned during the conduct of an investigation. There also may be times when your efforts and your methodologies in investigating, prior to involving law enforcement, will come under scrutiny. Probably the most disconcerting of all situations is where law enforcement is compelled to make public disclosure of an incident. It is during these times that you want to work with your own management and legal staff as well as with the law enforcement agency conducting the investigation. Everyone concerned wants to make the public aware of any known or suspected threats, but this must also be weighed against the potential harm that might result from a public release. Another concern from the company's

point of view is public perception of security measures within the organization, and to what extent they can protect information and thus customer investments.

61.7 CONCLUSION. Trying to stay ahead of cybercriminals is not getting easier, though inroads are being made, especially into the group Anonymous. Attempts to understand their intent or motivations are rarely successful. The best chance any organization has is to come to know local law enforcement agencies and to help them know you. Articulate your expectations in the event of a significant cyberevent, and have them do the same. Working together is the only path to success. Given that not all computer security or cyberlaw enforcement entities are created equal, you have to be willing to share knowledge, all in pursuit of the same end goal.

61.8 FURTHER READING

Ballou, S., and K. Higgins, eds. *Electronic Crime Scene Investigation: A Guide for First Responders*. National Criminal Justice Reference Service of the Office of Justice Programs, U.S. Department of Justice, 2001. Available free from www.ncjrs.gov/pdffiles1/nij/187736.pdf

Computer Crime and Intellectual Property Section. Criminal Division, United States Department of Justice. "Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations," 2002. Available free from www.cybercrime.gov/s&smanual2002.htm (URL inactive)

Web sites

InfraGard home page: www.infragard.org
www.fbi.gov/cyberinvest/cyberhome.htm
www.securityfocus.com/infocus/1523
www.ic3.gov/complaint/default.aspx

61.9 NOTES

1. Information provided by Jürgen Pabel, MSIA, Office for Security in Information Technology, Information Security Officer at Deutsche Post AG.
2. See www.ice.gov/cyber-crimes
3. See http://i.i.com.com/cnwk.1d/i/tr/downloads/home/computercrimereportingchecklist.pdf
4. www.usdoj.gov/criminal/cybercrime/cccases.html (URL inactive)
5. "Learn More About Infragard," Infragard Web site, December 9, 2011, www.infragard.org

INTRODUCTION TO PART VII

MANAGEMENT'S ROLE IN SECURITY

Management responsibilities include judgments of which resources can rationally be expended in defending against which threats. Managers must understand how to cope with the lack of quantitative risk estimates while using what information is available to guide investment decisions in personnel and technology. Their decisions are affected by regulatory and legal requirements and by the practical constraints of their relationships with other leaders within their organizations. This part includes chapters and topics that bear on information assurance managers' roles:

62. **Quantitative Risk Assessment and Risk Management.** Which vulnerabilities warrant repair? Which threats must be taken seriously? How much expense is justified on specific security measures?
63. **Management Responsibilities and Liabilities.** Roles, responsibilities, due diligence, staffing security functions, and the value of accreditation and education
64. **U.S. Legal and Regulatory Security Issues.** For U.S. practitioners especially, this chapter reviews the Gramm-Leach-Bliley Act and the Sarbanes-Oxley legislation
65. **The Role of the CISO.** The chief information security officer as an agent of change and as a strategist working to ensure that security fits into the strategic mission of the organization, and that it is communicated effectively to other C-level executives
66. **Developing Security Policies.** Approaches to creating a culture of security where policies grow organically from the commitment of all sectors of the organization, instead of being imposed unilaterally by security staff
67. **Developing Classification Policies for Data.** The essential role of data classification and how to implement systems that conform to regulatory and legal requirements
68. **Outsourcing and Security.** Security of outsourcing and outsourcing of security

CHAPTER 62

QUANTITATIVE RISK ASSESSMENT AND RISK MANAGEMENT

Robert V. Jacobson and Susan Baumes

62.1 AN INTRODUCTION TO RISK MANAGEMENT 62·2

62.1.1 What Is Risk? 62·2
62.1.2 What Is Risk Management? 62·2
62.1.3 Regulatory Compliance and Legal Issues 62·3

62.2 OBJECTIVE OF A RISK ASSESSMENT 62·4

62.2.1 Qualitative versus Quantitative Risk Assessments 62·4

62.3 LIMITATIONS OF QUESTIONNAIRES IN ASSESSING RISKS 62·6

62.4 A MODEL OF RISK 62·6

62.4.1 The Two Inconsequential Risk Classes 62·6
62.4.2 The Two Significant Risk Classes 62·7
62.4.3 Spectrum of Real-World Risks 62·7

62.5 RISK MITIGATION 62·10

62.5.1 ALE Estimates Alone Are Insufficient 62·10
62.5.2 What a Wise Risk Manager Tries to Do 62·10
62.5.3 How to Mitigate Infrequent Risks 62·13
62.5.4 The ROI-Based Selection Process 62·15
62.5.5 Risk-Assessment/ Risk-Management Summary 62·16

62.6 RISK-ASSESSMENT TECHNIQUES 62·16

62.6.1 Aggregating Threats and Loss Potentials 62·17
62.6.2 Basic Risk-Assessment Algorithms 62·17
62.6.3 Loss Potential 62·18
62.6.4 Risk Event Parameters 62·21
62.6.5 Threat Effect Factors, ALE, and SOL Estimates 62·22
62.6.6 Sensitivity Testing 62·23
62.6.7 Selecting Risk Mitigation Measures 62·24

62.7 SUMMARY 62·24

62.8 FURTHER READING 62·24

62.9 NOTES 62·25

62.1 AN INTRODUCTION TO RISK MANAGEMENT

62.1.1 What Is Risk? There is general agreement in the computer security community with the common dictionary definition of risk: *the possibility of suffering harm or loss*. The definition shows that there are two parts to risk: the *possibility* that a risk event will occur, and the harm or *loss* that results from occurrences of risk events. Consequently, the assessment of risk requires consideration of both factors: the frequency of threat events that cause losses, and the loss each such event causes. The product of the two factors, frequency and loss, is a *quantitative* measure of risk; for example, we might express a risk as an *Annualized Loss Expectancy* (discussed in Section 62.2) of, say, $100,000 per year. Risk is managed by taking actions that recognize and reduce these two factors.

Management of risk is important to the design, implementation, and operation of information technology (IT) systems because IT systems are increasingly an essential part of the operation of most organizations. As a result, both the size of the potential harm or loss and the possibility of a risk event occurring are increasing. In extreme cases, a risk loss may be large enough to destroy an organization. If an organization's risk-management program is deficient in some area, the organization may suffer excessive harm or loss, but at the same time the organization may be wasting resources on ineffective, excessive, or misdirected mitigation measures in other areas.

62.1.2 What Is Risk Management? This chapter discusses risk management with special emphasis on two aspects of risk management: risk assessment and risk mitigation. Exhibit 62.1 below suggests that risk management can be thought of as a four-step process.

Step 1: Identify and assess the risks. A risk assessment will only be as good as the inventory of systems and processes. If an organization does not have an adequate inventory, there may be systems or processes with significant risk that may not be reviewed. Once an inventory is complete, the organization can determine the risks associated with each.

Step 2: Identify the *optimum* set of risk-management measures.[1] In this context, *optimum* means making the best use of the resources available for risk

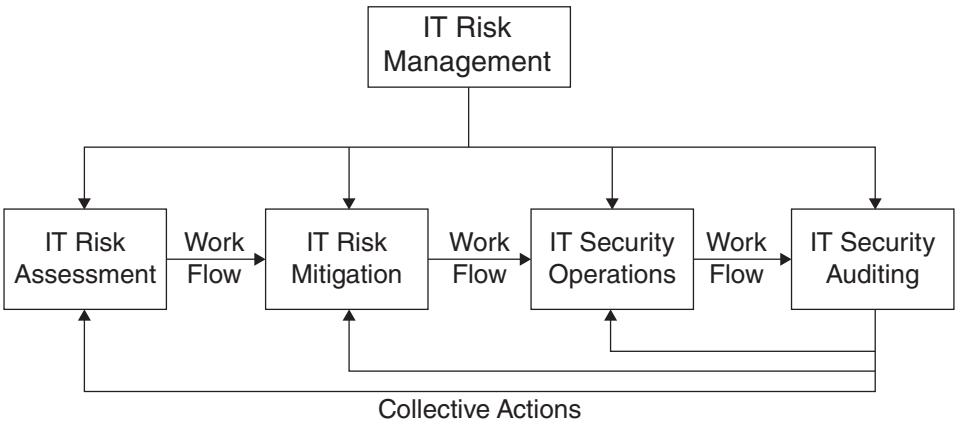


EXHIBIT 62.1 The Four IT Risk-Management Activities

management. Because no organization has infinite resources, risk management requires allocating finite resources by selecting from the broad range of IT risk mitigation measures, the *optimum* set. Since the resources being allocated are measured quantitatively in monetary terms, the risks must likewise be measured in monetary terms. A CIO can hardly expect to go to the CFO and announce that the risk assessment has identified a *high* risk, so a *high* security budget will be needed to address the risk. This issue is discussed in more detail in Section 62.4.3.

Step 3: Perform the ongoing, activities necessary for the effective functioning of the risk-management measures. Depending on the mitigating controls implemented, this could be ongoing log reviews of systems or quarterly meetings with local police. It will depend on the risk and the mitigating controls implemented.

Step 4: Security auditing is conducted to evaluate the effectiveness of IT security operations and to detect changing conditions that require a reassessment of risks, the implementation of new mitigation measures, or modifications to the security management program. Daily security management tasks and IT auditing are discussed elsewhere in this *Handbook* (see, for example, Chapters 23, 27, 40, 47, 52, 53, and 54), so this chapter will confine itself to the first two steps: *risk assessment* and *risk mitigation*.

62.1.3 Regulatory Compliance and Legal Issues. Increasing regulatory and legal issues are impacting risk management, most notably the Sarbanes-Oxley Act. More and more corporate officials are being sent to jail for infractions. As a result, it is not surprising that there may be an overreaction to compliance regulations, with excessive measures being recommended by auditors. Whenever implementation of a compliance recommendation does *not* include a cost justification, and the implementation cost is material, it may be a candidate for a simple risk analysis like this.

1. What effect does the recommended measure have on the probability of noncompliance? In other words, what will be the difference in the likely loss experience with and without the recommended measure?
2. Does this difference compare favorably with the cost to implement and maintain the recommended measure?

There are several points to be made.

- If the cost-benefit is marginal or negative, consideration should be given to alternate (and lower cost) ways to achieve the same end result.
- If the cost-benefit is strongly negative, consideration should be given to dropping the recommendation.
- Under no circumstances should an organization deliberately break the law.

For more about these issues, refer to Chapter 65.

Readers should note that there are strong disagreements about the validity of using quantitative risk-assessment and risk-management methods. For example, Donn Parker has argued for decades that quantitative estimates of risk are an illusion and that expert consensus on best practices provide a far more solid foundation for risk management.[2]

Chapter 10 in this *Handbook* includes discussion of the problems of ascertainment that complicate quantitative estimation of risk.

62.2 OBJECTIVE OF A RISK ASSESSMENT. This *Handbook* describes the pros and cons of a wide range of security measures that can be applied to an IT system. However, no organization has unlimited resources, and so it is necessary to decide which measures should be implemented.[3] It would be convenient if there were a set of security standards that applied equally to all IT systems, but as guidelines cited in Section 62.1.3 note, this is not the case for two basic reasons:

- Each IT system has its own particular risk environment.
- Each IT system has a unique workload. Although two or more IT systems may perform the same list of functions, it is highly unlikely that the systems will have exactly the same level of activity for each of the functions, and so the cost of lost data and service interruptions will not be the same.

Because of these differences, each IT system will have unique security requirements.[4] These two factors correspond to the two elements in the definition of risk:

- The risk environment determines the possibility of experiencing harm or loss.
- The characteristics of the system workload, and the associated IT system assets, determine the magnitude of the harm or loss.

The objective of a risk assessment is to generate information about risk exposures and the potential for loss associated with the workload in order to estimate the annualized loss expectancy (ALE) of the IT system under current conditions.

62.2.1 Qualitative versus Quantitative Risk Assessments. There are two basic types of risk assessments: *qualitative* and *quantitative*. The difference between them is substantial. A quantitative assessment provides a dollar figure associated with the loss. In the case of a qualitative risk assessment, the dollar figure may not be easy to obtain. The cost of the loss is difficult or in some cases impossible to quantify. As an example, if a retailer's database containing customer records is breached, how does one quantify the reputational impact in financial terms? The retailer may have significant losses associated with the breach, but for how long? Loyal customers may continue to frequent the business, but there may be a loss of new clients. Obviously, the impact will depend on the type of business and the relationship of the customer to the business.

In the example of the retailer performing a risk assessment, the customer impact may be rated as Low, Medium, or High with set boundaries. Oftentimes an organization will incorporate a combination of qualitative and quantitative risk-assessment properties in their final risk analysis. Certainly, organizational maturity plays a role in the type of risk assessment completed, particularly one that has never performed any type of risk analysis. In instances where no analysis has been completed, the firm may start with a qualitative analysis and then apply quantitative analysis to the top 10 risks to help prioritize a mitigation strategy.

A quantitative assessment provides a dollar figure associated with a loss. As organizations review their risk mitigation strategy, having the dollar figure is extremely

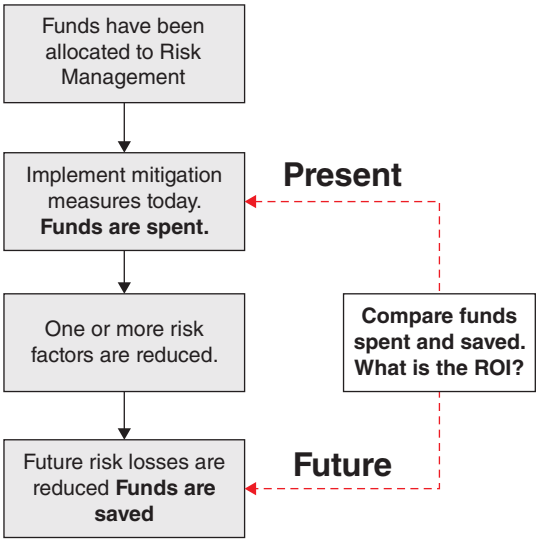


EXHIBIT 62.2 Evaluating an IT System Security Strategy

helpful. The dollar estimate is used to optimize the subsequent risk mitigation decisions and to validate decisions about existing security measures. In this context, *optimize* means to allocate the organization's resources to those actions that will yield the best overall performance. It would be suboptimum to spend $10,000 to avert $1,000 in expected losses. To avoid such wasteful expenditures, one must be able to estimate quantitatively both expected losses[5] and the effect of proposed mitigation measures. For this reason, in all but the simplest situations a risk assessment must produce a quantitative, monetary measure of risk, so that (a) risks can be compared with one another on a common basis and (b) the cost of risk mitigation measures can be related to the risks they are meant to address. Assessing a risk as *high*, *unacceptable*, or in other qualitative terms does not provide the information needed to support decisions to implement risk mitigation measures, which will always have quantitative implementation costs.

Exhibit 62.2 above shows why a quantitative estimate of ALE is an essential element in the optimization of the strategy. A sum of money has been budgeted for IT risk management for next year. A risk-management measure is selected and implemented, and so funds are expended. The implementation is expected to reduce one or more risk factors. As a result, we expect *future* risk losses will be reduced, and funds will be saved. By comparing the present values of the funds expended now and those saved in the future, we can estimate the return on investment (ROI) of the implemented measure. The goal is to maximize the ROI, and to avoid risk-management measures with small or negative ROIs. This is the essential character of businesslike IT system risk management. However, it is evident that we must be able to make credible estimates of expected risk losses under both conditions, and under the assumption that a proposed risk-management measure has been implemented to optimize our risk-management strategy.

It may be suggested that quantitative assessment of risk is *too complicated*, and so a simpler, qualitative method can be used instead. However, the complexity is inherent in risk itself, and any *simple* assessment technique will be inherently weak, and potentially misleading. This topic is discussed in more detail in Section 62.4.3.

62.3 LIMITATIONS OF QUESTIONNAIRES IN ASSESSING RISKS. One may be tempted to think of a *security questionnaire* or checklist as a valid risk assessment, but even if the checklist seems to be authoritative because it is automated, this is not the case. Typical questionnaires can help identify inventory and therefore potential areas of risk, but *not* expected loss. At best, a questionnaire can only compare the target of the questionnaire with a security standard implicit in the individual questions. Questionnaires do not meet the objective of a risk assessment, because the answers to a questionnaire will not support optimized selection of risk mitigation measures. The most a well-designed questionnaire can do is to identify potential risk areas, and so they can be of help in scoping and focusing a quantitative risk assessment.

Questionnaires suffer from several other shortcomings:

- The author of a question and the person who is answering the question may have different understandings of the meanings of key words. As a result, an answer may not match the intent of a question.
- Questions tend to be binary, but most answers are inherently quantitative. For example, consider the binary question "Do users comply with password policy?" Although the anticipated answer is either yes or no, a meaningful answer is likely to be much more complicated because each user will have a unique pattern of compliance. Some users will make every effort to be 100% compliant, some will make every effort to circumvent the policy, and the remaining users will comply to a greater or lesser extent. Note also that the question does not attempt to evaluate the appropriateness of the policy.
- Because questionnaires are inherently *open loop*, and because of the binary answer issue discussed above, questionnaires tend to miss important information that would be elicited by a risk analyst who interacts directly with respondents and responds to a quantitative answer with a new question.[6]

Questionnaires are appealing because they appear to relieve the risk assessor of the need to probe into the business system being assessed, but something better is needed for a valid risk assessment. The ensuing sections of this chapter will build a model of risk, and then show how to apply quantitative parameters to the model. Although a greater effort is required, far superior results are achieved.

62.4 A MODEL OF RISK. Exhibit 62.3 is a simplified model of threats and consequences devised by Robert Jacobson. The model takes the first steps toward quantification of risks. In this very simple model, all risk events are assumed to have either a low or high rate of occurrence, and all consequences of risk event impacts are assumed to be either low or high. Mr. William H. Murray, at the time an executive consultant at Deloitte & Touche, referred to this risk model as Jacobson's Window. Let us now consider the implications of the model.

62.4.1 The Two Inconsequential Risk Classes. The two-by-two empty matrix model implies that there are four classes of risk: low-low, high-low, low-high, and high-high. Exhibit 62.4 suggests that two of the classes can be ignored. The low-low class can be ignored because these risks do not matter. As an extreme example, it is obvious that a risk event that occurs at about 10,000-year intervals and causes a $1 loss each time can be ignored safely. Experience suggests that the high-high class can be assumed not to exist in the real world. If 50-ton meteorites crashed through the roofs

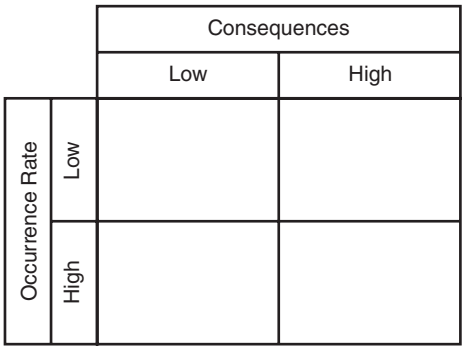


EXHIBIT 62.3 Jacobson's Window, A Simple Risk Model

of computer rooms every day, there would be no attempt to use computers. In practice, high-probability, high-loss risks just do not exist. Catastrophic events do occur, but not frequently.

62.4.2 The Two Significant Risk Classes. This analysis suggests that there are only two significant risk classes: high-low and low-high. Data entry keystroke errors are an example of a high-low risk: a high probability of occurring, and usually a low resulting loss. A major fire that destroys the building housing an IT system is an example of a low-high risk: a low probability of occurrence, and a high consequential loss. However, we know that real-world risks do not fall into just these two classes. Instead there is a spectrum of risks from high-low to low-high.

62.4.3 Spectrum of Real-World Risks. Exhibit 62.5 illustrates the distribution of representative risks from high-low (key stroke errors) to low-high (major fires). Conceptually, there is little difference between the high-low and low-high threats. Experience suggests that averaged over the long term the high-low and low-high threats will cause losses of similar magnitude to an organization. This concept is quantified by the notion of annualized loss expectancy (ALE). The ALE of a risk is simply the product of its rate of occurrence, expressed as occurrences per year, and the loss resulting from a single occurrence expressed in monetary terms, for example, dollars per year.

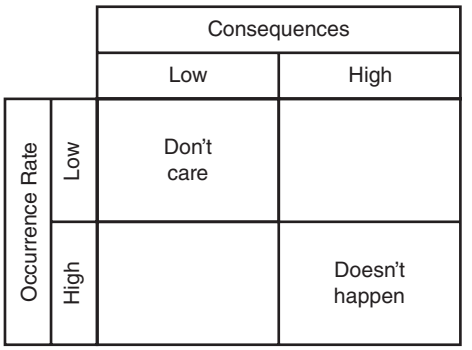


EXHIBIT 62.4 The Two Inconsequential Risk Classes

		Consequences	
		Low	High
Occurrence Rate	Low		"low-high" major fire flooding cash fraud
	High	power failure software bug key error "high-low"	

EXHIBIT 62.5 The Spectrum of Real-World Risks

Here is a simple example of ALE: Assuming that 100 terminal operators each work 2,000 hours per year and make 10 keystroke errors per hour, the occurrence rate of keystroke errors would be 2,000,000 per year. Next, assuming that 99.9% of the errors are immediately detected and corrected at an insignificant cost, 2,000 errors slip by each year (a high occurrence rate risk), which must be corrected later at a cost estimated to be $10 each, a low consequence for each occurrence. Thus, the ALE of the high-low risk: Keystroke error is estimated to be 2,000 occurrences per year × $10 per occurrence, or an ALE of $20,000/year.

Continuing our example, assume that the probability of a major fire in any one year is 1/10,000 (in other words the annualized probability of occurrence of a major fire is estimated to be 0.0001), and if the loss resulting from a major fire would be $200 million, then major fire is a low-high risk. These assumptions lead to an estimate of $20,000/year for the ALE of major fires.

Thus, the two risks from opposite ends of the risk spectrum are seen to have ALEs of about the same magnitude. Of course, we have manipulated the two sets of assumptions to yield similar ALEs, but the results are typical of the real world. However, note that we are assured of having a loss of $20,000 *every year* from keystroke errors, but a loss of $200 million in one of the next 10,000 years.

Exhibit 62.6 is a computer-generated[7] plot of threat occurrence rates and consequences (total single occurrence loss) taken from an actual quantitative risk analysis of a facility. Each of the numbered boxes represents a threat type. Because the two scales are logarithmic, the contours of constant ALE are diagonal straight lines. As expected there are no threats in the high-high (Doesn't Happen) and low-low (Trivial) zones.

While the ALEs span a range of about four orders of magnitude (from about $50 per year to about $500,000 per year), frequencies and consequences each span about seven orders of magnitude. This result is typical of the real world, and illustrates the fundamental flaw in High-Medium-Low risk models like the matrix in Exhibit 62.7. These models attempt to simplify risk assessments by defining a very limited range of values, three each, for frequency and consequence. As a result, threat event ALEs that differ by as much as two orders of magnitude, as illustrated in Exhibit 62.5, are given the same ranking when assessed using a high-medium-low matrix.

Proponents of this risk model are silent about the details. Firstly, as Exhibit 62.6 shows, the *Very High* and *Very Low* risks are meaningless in the real world. Secondly, the matrix requires users to compress numerical values with a range of almost seven

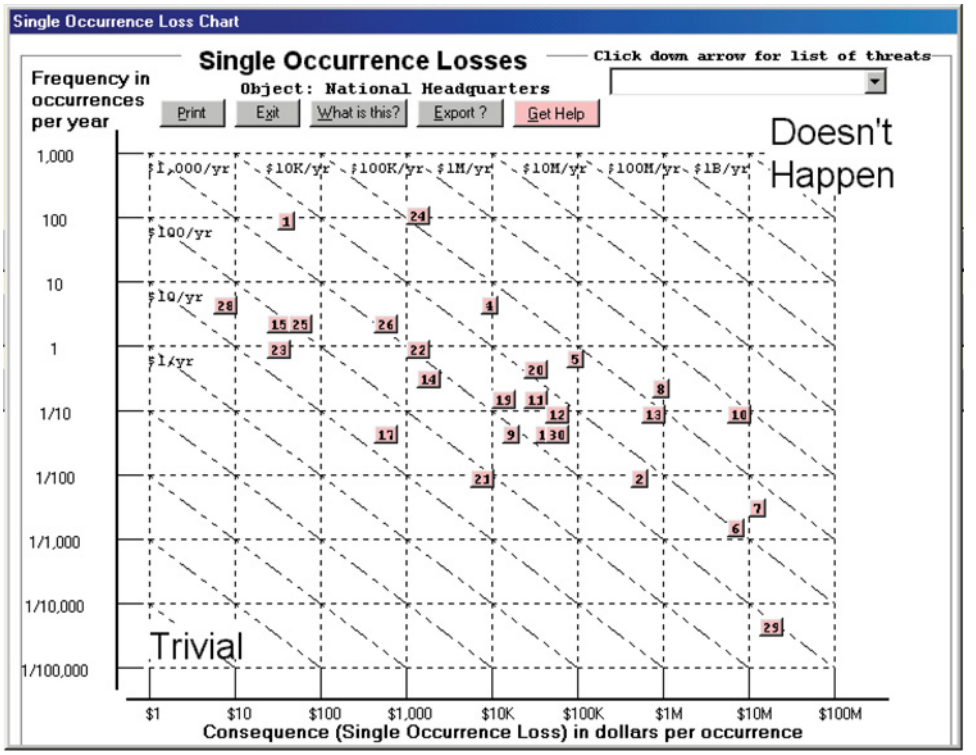


Exhibit 62.6 A Plot of Actual Threat Frequencies and Consequences
© Copyright 2005 International Security Technology, Inc.

orders of magnitude, for example, from 100/year to 1/20,000 years, into only three values: Low, Medium, and High. In other words, two risks with same consequence but occurrence rates that differ by a factor of 100 might receive the same evaluation. Likewise, because of the arbitrary boundaries between low, medium, and high, two risks with the same consequence but with occurrence rates that differ by 1% could receive different ratings. These considerations defy common sense, and make it clear

		Threat Occ. Rate		
		Low	Medium	High
System Impact	High	Medium	High	Very High
	Med.	Low	Medium	High
	Low	Very Low	Low	Medium

Exhibit 62.7 A Typical Risk Matrix

that rating a risk as High, Medium, or Low cannot be the basis to cost-justify a proposed risk management measure.

62.5 RISK MITIGATION. Exhibit 62.6 shows that the effect of risk exposures on IT systems can range from trivial to catastrophic, and it is not always immediately obvious which risk exposures are the most dangerous. For this reason, it is essential to base the selection of risk mitigation measures on a quantitative assessment of risks. In this section we consider the practical considerations in generating quantitative assessments and applying them to risk mitigation decisions.

62.5.1 ALE Estimates Alone Are Insufficient. As noted above, ALE is a useful concept for comparing risks, but we recognize intuitively that ALE alone is not a sufficient basis for making risk mitigation decisions about the low-probability high-consequence risks. There are two reasons for this.

- The first reason is the difficulty in generating a credible estimate of occurrence rate for low-probability risks. As a rule, one can generate credible estimates of the consequences of a low-probability risk, but the same is not true of its occurrence rate. Risks that flow from human actions such as fraud, theft, and sabotage are also difficult to quantify credibly.
- The second reason stems from what appears to be a common human trait that this writer has postulated as Jacobson's 30-Year Law:

People (including risk managers) tend to dismiss risks that they have not *personally* experienced within the past 30 years.

Why 30 years? It is not clear, but it may be related genetically to human life expectancy, which until just a few generations ago was about 30 years. Possibly, people who were able to suppress anxiety about rare events were more successful than those who worried too much. Numerous instances of Jacobson's 30-Year Law can be found. For example, the United States Government has had a major fire at about 28-year intervals beginning in 1790, most recently at the Military Records Center. Presumably each new generation of federal property managers must relearn the lessons of fire safety by direct experience. The Northeast power blackout of 2003 followed a similar event in 1976, 27 years earlier.

It seems to be common for senior managers, particularly public officials, responding to a calamity such as the meltdown of the Fukushima nuclear plant in March 2011 following an earthquake and the resulting tsunami, to say something like this: "Who could have imagined that such a thing would happen? However, we have taken steps to see that it will never happen again."[8] This is an imprudent statement for two reasons. First, prudent managers will anticipate potential disasters. Second, saying *never* implies that perfect security is the goal. This is nonsensical, since perfect security is infinitely expensive, and so cannot be achieved.

62.5.2 What a Wise Risk Manager Tries to Do. Unlike the senior managers quoted above, an organization's risk manager should be trying to imagine every possible *material* risk the organization faces, even those not personally experienced, and developing estimates of the impact of these risks. Next, the risk manager should strive to identify the optimum response to each material risk by identifying security measures that have a positive return on investment (ROI). Note that potentially fatal

low-occurrence/high-consequence threats require treatment irrespective of ALE and ROI as discussed in Section 62.5.3.

Consider this example of cost effectiveness (presented as individual steps in reasoning):

- The ALE for keystroke errors was estimated in the Section 62.4.3 example as $20,000 per year, or $200 per operator.
- This estimate is credible because presumably there has been ample past experience with both the occurrence rate and impact cost of the keystroke-error risk.
- The risk manager considers how to treat this risk.
- Imagine that experience suggests that spending $100 each year on keyboard skills training for each operator would reduce the undetected error rate by 30%.
- Because the $100 per operator expense would yield a benefit of $60 in reduced ALE (30% of $200, the per-operator ALE), the ROI[9] of training would be –40%, clearly not a cost-beneficial mitigation measure.
- In effect we would be spending $10,000 to achieve a $6,000 loss reduction.
- However, if the error rate were reduced by 90% instead of 30%, the training would appear to be a good investment; spending $100 each to train operators, would produce a reduction in ALE per operator of $180, an ROI of +80%.
- The goal of the risk manager is to find the package of risk-management measures, which yields the greatest overall ROI.

62.5.2.1 The Three Risk-Management Regions. Basing risk-management decisions on ROI estimates, as described above, does not work for the low-probability, high-consequence risks because of two negative factors (1) the credibility of risk estimates and (2) management's concern about the risk consequences. End users commonly have a higher level of concern about risks than IT system managers, but as a rule it is the IT system managers who make the decisions about security measures. An IT system manager generally has no difficulty choosing between buying a faster server, and a more reliable server. The benefit of higher throughput is immediately evident. The benefit of higher reliability is not as obvious. Thus, although a risk manager may identify a significant low-occurrence, high-cost risk before it has occurred, an organization's senior managers probably would be unaware of these low-high risks, and would be genuinely surprised were a major loss to occur. Hence, the "Who could have imagined . . ." press releases. In other words, simply identifying an exposure might be thought to be enough to justify adoption of a mitigation measure to address the exposure. This suggests that the risk manager needs additional criteria for selecting risk mitigation measures. The next section presents an overview of mitigation measures to clarify the selection criteria.

It will help risk managers to understand the universe of risk events to which an organization is exposed, by dividing the total risk space into the three regions shown in Exhibit 62.8. The two axes are logarithmic, and the dots represent individual threats.

These regions help define appropriate risk-management actions. The bounders of the regions are defined by two senior management decisions.

1. The first decision is to define the *minimum significant threat event occurrence rate*. The concept is that it is reasonable to simply ignore the risk of threat events for which we have estimated occurrence rates less than some minimum rate. For

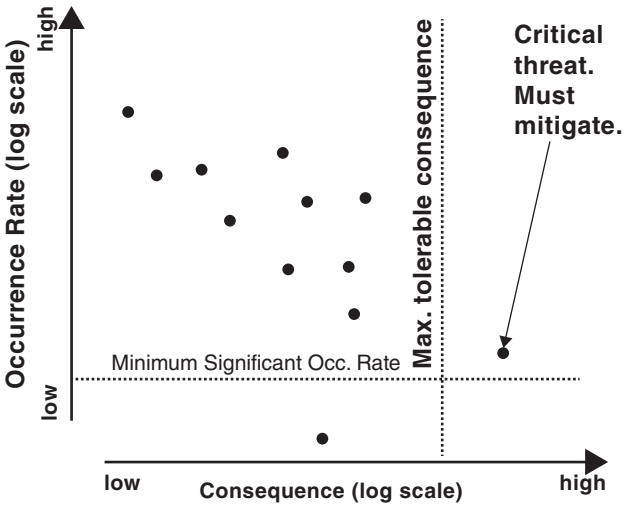


EXHIBIT 62.8 The Three Risk Zones
© Copyright 2002 International Security Technology, Inc.

example, senior management may decide to ignore risks with occurrence rates estimated to be less than once in 20,000 years—that is, the probability of an occurrence of the threat next year is estimated to be less than 0.00005.

2. Senior management may also identify a loss level (consequence) that is *intolerably high*. Risk events of this type will appear in the upper right of Exhibit 62.6. The occurrence rate of these events is immaterial as long as they exceed the minimum consequence criterion and the minimum occurrence rate. If we estimate that the loss caused by an occurrence of a threat event (commonly referred to as the *single occurrence loss*, or SOL) exceeds the loss threshold *and* the occurrence rate exceeds the minimum material occurrence rate, then we *must* take steps to reduce the loss, perhaps by transferring the risk with an insurance policy, or by reducing the estimated occurrence rate to a value below the minimum material occurrence rate.

It is instructive to consider where to plot an unusual risk event like the attack on the World Trade Center on September 11, 2001. It is likely that many organizations have facilities that would generate losses in excess of their maximum tolerable loss if totally destroyed by a similar terrorist attack. This implies that such organizations should take steps to protect these facilities against terrorist attacks. However, we should also take into account the estimated rate of occurrence before making a decision. How can we make a credible estimate of such a rare event? We can begin by estimating how many such events will occur next year in the United States. For example, based on past experience we might assume that there will be about two such attacks each year. Secondly, we estimate the likelihood that, of all the facilities in the United States, our facility would be selected for attack. This will depend on the *attractiveness* of our facility to a terrorist group when compared with all other potential targets in the United States. For example, assume that there are 100 *attractive* buildings that are as widely recognized and as vulnerable as the WTC. If our facility is one of the hundred, then we might estimate the occurrence rate as $1/100 \times 2$, or a probable occurrence rate of

0.02/year. On the other hand, if our building is *not* one of the 100 high-profile buildings, then we might estimate that if there are 200,000 similar buildings in the United States, our occurrence rate is 1/200,000 × 2 or an occurrence rate of 1/100,000 years. In this latter case, our senior management probably would choose to ignore the risk.

The remaining portion of the Exhibit 62.6 graph, the mid- to upper-left-hand corner, is the zone where the risk manager attempts to find cost/beneficial risk-management measures as discussed in the sections that follow.

62.5.2.2 Where ROI-Based Risk Management Is Effective. ROI-based mitigation works well for high-probability, low-consequence risk exposures for two reasons: The manager who approves the expenditure believes that the risk exists and should be addressed, and believes further that the parameters used to generate estimates of ALE and the reduction in ALE, used to estimate ROI, are reasonable and credible.

62.5.2.3 Four Reasons for Adopting a Risk-Management Measure. There are four general tests of the utility of a risk-management measure:

1. The measure is required by law or regulation. In effect, a governing body has determined (one hopes) that the measure makes good public policy because it will always meet one of the remaining three tests. Exit door signs in public buildings are a good example. Sarbanes-Oxley may be an example of a bad (not cost-effective) law.
2. The cost of the measure is trivial, but its benefit is material. For example, a little-used door, which compromises physical access controls, is not being kept locked. One can institute a procedure and install security hardware to keep the door locked at very low cost.
3. The measure addresses a low-high risk that has an intolerable SOL, as discussed in Section 62.5.2.1. For example, it would be intolerable for a corporation to experience an SOL that exceeded owner equity or net worth. The failure several years ago of a prominent British merchant bank following unwise speculation by a staff member is a tragic example of an organization that failed to identify and address an intolerable SOL exposure.
4. The cost of the measure will be more than offset by the reduction in future losses (ALE) that it will yield. In other words, the mitigation measure has a positive ROI. This reason is commonly used to justify protection against the high-low risks. Operator keyboard training described in Section 62.5.2 is an example.

Procedures for managing the last two risk categories follow.

62.5.3 How to Mitigate Infrequent Risks. After the high-low threats have been addressed using an ROI analysis, the risk manager considers all imaginable low-high risks, one by one, and makes for each such risk an estimate of the SOL and the rate of occurrence. The report of this analysis should describe for each threat the confidence level of the estimates of its SOL and occurrence rate. These estimates are arranged in descending order of SOL. The list is presented to senior management, who draws a line somewhere on the list and says, "The risks above the line are intolerably high. Do something about them." The risk manager then considers each of the unacceptable risks in two ways.

62.5.3.1 Reduce the Magnitude of High Single-Occurrence Losses.

Sometimes the magnitude of an SOL can be reduced. There are several possibilities:

1. Transfer the risk by obtaining insurance against it. The premium will depend in part on the amount of the loss that is deductible. For example, one might obtain insurance against a $100 million SOL with a $10 million deductibility to minimize the insurance premium. In effect, the intolerable $100 million SOL has been reduced to a tolerable $10 million SOL at the cost of the insurance policy premium.
2. Disburse the risk exposure. For example, replace a single IT center with an intolerable SOL of $500 million of catastrophic physical damage, and service interruption losses, with three centers having SOLs of about $167 million each. The centers should be sufficiently isolated from one another to rule out shared disasters. The cost will be the incremental cost of the less efficient operation of three facilities.
3. Reduce the vulnerability of the IT system to the risk. For example, implementing an enhanced business resumption plan, at some additional cost, can speed up recovery off site. This will reduce the SOL associated with catastrophic service interruption losses. Notice that this an example of a risk mitigation measure that affects more than one threat.

The risk manager also may strive to reduce the occurrence rate of a high SOL. Because of the uncertainty of the estimates of low occurrence rates, this is less satisfactory. Nonetheless, even the uncertain occurrence rate estimates can be useful. If two risk exposures have the same SOL, but differ by an order of magnitude in estimated occurrence rate, it is reasonable to assume that the risk with the lower occurrence rate represents a lesser danger to the organization.

62.5.3.2 Risk-Management Measures Selection Process.

Assume that the risk manager has presented a risk-assessment/risk mitigation report to senior management as described above. The report lists the low-high risks with one or more strategies for treating each risk.[10] The senior manager is responsible for selecting the mitigation measures for implementation because, as noted above, ROI-based justification is inappropriate for this class of risks.

1. If the risk manager is able to identify a relatively low-cost mitigation measure that reduces the rate of occurrence of a risk to a low enough value, senior management may elect to adopt the mitigation measure.
2. There is some rate of occurrence below which the senior manager is willing to ignore a risk, even if the estimate is low confidence. Typically, this applies to very low occurrence rate events. Extreme examples are a nuclear detonation or a crashing meteorite.
3. Risk transfer by insurance, or one of the other techniques listed above, reduces the SOL to a tolerable level, and the senior manager is willing to accept the cost of the insurance coverage.

A complete tabulation of occurrence rates, and the costs of mitigating actions for all the high-SOL risk exposures, will help senior management to prioritize the implementation of the mitigation measures.

62.5.4 The ROI-Based Selection Process. After the potentially fatal risks have been addressed as described in the preceding sections, the risk manager can turn to selection of measures to address the remaining high-low risks. A five-step procedure for selecting strategies on the basis of ROI is straightforward.

1. For each of the threats, identify possible risk-management measures. As noted above, some measures will affect more than one threat. There are four possible strategies:
   - *Threat mitigation*. Measures to reduce the occurrence rate and/or the impact of one or more of the threats.
   - *Risk transfer*. Use an insurance policy or other contractual means to transfer some or all of the risk to another party.
   - *Business resumption plans*. Take steps in advance of threat occurrences to decrease the time required to resume business operations. This type of strategy applies to non-IT business systems that are included in the scope of the risk-management project.
   - *IT system recovery plans*. Take steps in advance of threat occurrences to decrease the time required to resume operation of the IT system, and to reduce the amount of stored data lost. See Section 62.6.3.2 for details.
2. For each strategy, estimate the present value of the ***cost*** to implement and maintain the measure, taking into account the remaining life of the facility, the useful life of the measure, and the applicable time value of money. The organization will not enjoy the full benefit of a measure with a high initial cost and a low annual support cost if the life of the facility is less than the life of the measure.
3. For each strategy, estimate the effect of the measure on the occurrence rate and consequence of the threat(s) that the strategy addresses. Use these data to calculate total ALE, assuming the measure is implemented. Calculate the present value of the ***reduction in total ALE*** the strategy is expected to yield. This is the ***benefit*** of the measure. Note also that in addition to reducing risks a strategy may reduce operational costs. This cost saving should be added to the ongoing benefit (ALE reduction) of the strategy.
4. Using the values of cost and benefit, calculate ROI as follows:

   $$\text{ROI} = ((\$\text{Benefit} - \$\text{Cost})/\$\text{Cost}) \times 100\%$$

   For example, using the data in Section 62.5.2, ROI is calculated as follows:

   ```
   ROI = (($6,000 - $10,000)/$1,000) × 100%
       = (-$4,000/$1,000) × 100%
       = 0.4 × 100% = -40%
   ```

   This equation normalizes the ROI and expresses it as a percentage. Zero percent represents a neutral ROI (benefit equals cost). One hundred percent would represent a very beneficial measure, and 1,000% would probably be a no-brainer. If the measure generates a cost saving, the saving is added to the ALE reduction benefit when calculating the ROI. *Note that the ROI analysis of a risk mitigation measure* ***will understate its benefit****, if the risk assessment does not include all the risk events that the measure will address.*

5. Make a list of the strategies in descending order of ROI after discarding the strategies with a negative ROI.

Presumably, management now selects for implementation the strategies in descending order of ROI until all available resources are allocated. However, if two or more strategies impact the same threat, it is possible that their effects may be partially overlapping. The risk manager can investigate this possibility the following way. Begin by assuming that the most attractive strategy has been implemented, and recalculate the *baseline ALE*. Next, verify the parameters for the remaining strategies, taking into account the effect of the measures assumed to have been implemented. Recalculate their ROIs as described above, and repeat the selection process. It would be convenient if there were an analytic way to do this, but this author has been unable to devise a practical way to do so because of the wide range and complexity of potential interactions.

62.5.5 Risk-Assessment/Risk-Management Summary. The risk model illustrated by Jacobson's Window leads to the following conclusions:

1. Risks can be broadly classified as ranging from high-probability, low-consequence to low-probability, high-consequence events. In general, risk events in this range cause losses over a relatively narrow range of loss values when expressed as an annual rate or ALE.
2. Because of IT operating staff familiarity with the high end of the high-low risks, all available measures may already have been taken such that there may not be any cost-beneficial way to further reduce the ALE of some of these threats.
3. Midrange risks can be addressed by selecting mitigation measures with a positive ROI, based on the relationship between the cost to implement a mitigation measure and the reduction in ALE it is expected to yield.
4. Treatment of low-high risks requires the judgment of senior management, based on estimates of SOL and, to a lesser extent, estimates of rate of occurrence. Mitigation measures or risk transfer may reduce SOLs to acceptable levels or decrease occurrence rates to the level at which risks can be ignored.
5. To be effective, the risk-management function must be:
   - Performed by properly qualified persons
   - Independent of IT line management
   - Reported to senior management to ensure that all risks are recognized, and that resource allocation is unbiased

62.6 RISK-ASSESSMENT TECHNIQUES. The first step in performing a risk assessment is to define the scope and focus of the assessment. The scope will define the IT functions or business processes, the physical assets, the intangible assets, and the liability exposures, which are to be included in the assessment. The focus means the kinds of risk events to be included in the analysis. Risk events may include failures of the IT hardware and software, software logical errors, deliberate destructive acts by both insiders and outsiders, and external events such as earthquakes, hurricanes, river floods, and so forth. Sometimes the focus is determined by the desire to evaluate a specific class of risk-management measures.

62.6.1 Aggregating Threats and Loss Potentials. Aggregation is a key concept of quantitative modeling, and refers to the practice of combining related items to keep the model at a manageable size. For example, one can define a single risk event (*IT Server Failure*) or two risk events (*Server01 Failure* and *Server02 Failure*) under the assumption that the two servers support different systems with different loss potentials) and so on in ever-increasing detail. The more risk events defined for a risk assessment, the greater will be the effort required to establish the risk parameters, and perform the calculations. However, *the individual risk events will not all have the same impact on the results of the assessment*. Indeed, some risk events will be found to be completely inconsequential, so the time expended to estimate their parameters will have been wasted. In other words, an appropriate degree of detail will not be immediately obvious.

This suggests that one should begin an assessment using highly aggregated risks, functions (typically business processes), and assets, perhaps no more than 10 to 15 of each in a typical situation, together with roughly prepared estimates of the risk parameters. This preliminary analysis will show which elements are important to the results. One can then add details to the important elements as appropriate, and refine the material estimates to improve the validity of the model. This approach works to concentrate the effort on the important issues, and avoids wasting time refining unimportant input data.

62.6.2 Basic Risk-Assessment Algorithms. It is generally accepted that ALE, in dollars (or other currency as appropriate) per year, is calculated as follows:

ALE = Threat Occurrence Rate (number per year)
× Threat Effect Factor (0.0 to 1.0)
× Loss Potential (in monetary units)

The term *threat occurrence rate* (TOR) is used to designate the estimate of the probability of occurrence of a threat event during a risk-management planning period, typically one year, thus the TOR is the probability of occurrence at an annualized *rate*. The functions and assets are characterized by their potential to trigger a loss when impacted by one or more of the threats. The term *potential for loss* or simply *loss potential* is used specifically to designate the **worst-case loss** a function or asset can generate. This implies that there is at least one threat event (the worst-case threat) that will trigger a loss equal to the loss potential. Other threats will cause lesser losses because they have a lesser effect on the function or asset. The *Threat Effect Factor* is used to quantify the relationship between a threat event and a function or asset. If the threat event has no impact on the function or asset, the corresponding Threat Effect Factor will be zero. If the threat event triggers the worst-case loss (the loss potential), the Threat Effect Factor will be 1.0. Threat Effect Factors will range from zero to 1.0 for all threats.

Exhibit 62.9 illustrates how each threat-function or threat-asset pair includes a Threat Effect Factor. The SOL, expressed in monetary units for each function or asset with respect to a given threat event, is calculated as follows:

```
SOL = Threat Effect Factor (0.0 to 1.0)
      × Loss Potential (monetary units)
```

	Process #1	Process #2	Process #3	Process #4
Threat #1	TEF_{11}	TEF_{12}	TEF_{13}	TEF_{14}
Threat #2	TEF_{21}	TEF_{22}	TEF_{23}	TEF_{24}
Threat #3	TEF_{31}	TEF_{32}	TEF_{33}	TEF_{34}
Threat #4	TEF_{41}	TEF_{42}	TEF_{43}	TEF_{44}

EXHIBIT 62.9 Threat Effect Factors Connect Threats and Processes, Functions, and Assets

The terms *loss potential* and *Threat Effect Factor* are as defined above. The sum of the SOLs for all the functions and assets with respect to a given threat is the consequence of an occurrence of the threat.

From these two definitions, we can see that:

```
ALE = SOL × Threat Occurrence Rate ($/year)
```

The sections that follow discuss these two equations in more detail and describe techniques for estimating values for TOR, loss potential, and Threat Effect Factor.

62.6.3 Loss Potential. There are four basic kinds of losses that contribute loss potential:

1. *Property damage losses.* Property damage losses occur when a threat event impacts an asset of the organization. The asset may be physical property damaged, for example, by a fire or flood, or the asset may be an intangible asset, such as a trade secret or a proprietary database, the improper disclosure, modification, or destruction of which causes a loss to the organization.
2. *Liability losses.* The operation of an IT system may expose the organization to liability for damage or injury. For example, improper operation of an IT-controlled process might release a toxic gas. Improper disclosure of personal information may cause the individual to sue the organization for damages. If third parties place reliance on data generated or maintained by an IT system, there may be an exposure to damage suits if the data are incorrect as a result, for example, of flaws in the IT system.
3. *Service interruption losses.* Service interruption losses occur when IT system services or other business functions are interrupted, or are not initiated in a timely manner because of the action of a threat event. In general, the longer the duration of an interruption, the greater will be the amount of the loss. This means that service interruption loss potential must be estimated in the form of a table of values appropriate to the expected range of interruption durations the threats are expected to generate. As a consequence, each threat that causes

service interruptions must be defined in terms of the percentage of the threat's occurrences that result in each of the interruption durations used to define service interruption loss potential.

4. *Lost data losses*. IT systems typically require stored data to function. If the primary copy of the stored data is lost, the data are recovered from the most recent backup copy, or reconstructed from paper records, if they exist, of prior transactions. The magnitude of the data loss will depend on the rate at which the stored data change, the frequency with which a backup copy is made, the accuracy of the backup copy process, and the frequency of the occurrences of the threat events that require data recovery of the stored data. The stored data used by some IT systems do not change continuously, and so a complete backup copy can be made after each update phase, and data loss thus avoided.

Loss potentials, as described in the three following subsections, are based on estimates of the worst-case losses with respect to the list of risk events.

62.6.3.1 Property and Liability Losses. These losses are treated the same way, so they can be discussed together. It is advisable to begin consideration of property damage losses by making an aggregated list of possible sources. For example, the initial list might consider together all the hardware in an IT facility, and similarly all of the databases needed to operate each of the principal IT applications. For each such asset[11], an estimate is made of the worst-case loss that the organization could experience with respect to that asset. For example, the worst-case loss for a set of stored data would be the cost to reconstruct the data from back-up copies, and the cost to recover any missing or corrupted data. There will be other losses resulting from interruptions to IT services that require the availability of the data, but these losses are described separately in the next subsection.

Liability[12] loss potentials, if included in the scope of the risk-management project, should be estimated in the same way. The first step is to determine which liabilities should be included in the risk assessment. The organization may already have a mechanism for addressing liability exposure, possibly through a risk transfer mechanism. However, the liability exposures that are unique to the IT operations should be included, and the worst-case loss potential estimated as described above.

62.6.3.2 Service Interruption Losses. Service interruption (sometimes referred to as *denial of service*[13]) losses refer to the losses experienced by end-users of an IT system when the service is not performed on time. The first step in estimating service interruption losses is to construct a list of the IT services for which there is a significant interruption loss potential. Here, too, it is desirable to aggregate individual applications in basic service areas. Examination of a wide range of functions suggests that there are six different classes of service interruption losses. Since each class of loss may begin after a different interruption duration, it is unlikely that there will be a single loss rate that begins at the instant the interruption begins, and has a fixed rate per hour of interruption duration into the indefinite future. Reviewing the causes with the persons most familiar with each service, typically the line-of-business managers, will determine which classes apply and what the loss potential factors are. By tabulating the losses associated with each of the classes, one can generate an overall loss potential for a suitable range of service interruption durations.

- ***Reduced productivity***. Would people be idled if there were a service interruption? If yes, how many and how soon? What, approximately, is their total pay, including benefits? If substantial production facilities would be idled, what is the approximate hourly cost of ownership of the facilities, computed as the total life-cycle cost of ownership divided by the total operating life in hours?
- ***Delayed collection of funds***. Does the IT system application trigger collection of funds for the organization, for example, a billing system, or a loan-notes-due collection system? If yes, determine the average amount collected each business day, which would be affected by a service interruption. Because in most cases it will be possible to do some manual processing of the high-value items, the affected amount may be less than the full daily processing dollar volume. Determine a suitable cost-of-money percentage for use in present-value calculations by consulting the organization's treasurer or controller. Alternatively, the current commercial bank prime rate plus one or two percentage points can be used. The total loss does not increase linearly with the duration of the interruption. For example, assume that the one-day loss is $100 due to delayed collection. After a two-day outage, the loss would be $300: $200 from the first day and $100 from the second day. A three-day outage would cost $600. $300 of accumulated losses for the first day, $200 for the second day of delayed collections, and $100 for the third day, and so on.
- ***Reduced income***. Will an interruption of the application impact sales revenue or other income receipts? If yes, estimate the amount of income lost both immediately, and over the long term, for a range of service interruption durations appropriate to the risk events included in the risk assessment. It may be difficult for the end user to make an estimate, unless it is possible to remember the last major service interruption, and to describe what happened. What would the organization do if a major competitor had a service outage? How would the outage become known, and how would it be exploited to increase income? These questions may help to determine the amount of reduced income to be expected.

  Along with reduced income, there would probably be a reduction in operating expenses, such as cost of sales, caused by the decreased activity. When the reduction in operating costs is subtracted from the reduced revenues, the result is the net loss. Estimates should be made for the shortest service interruption duration that the user believes would cause significant lost business; in addition, planning should include provisions for one or more service interruptions of greater duration.
- ***Extra expense***. How would the user respond to an outage? Would it be necessary to hire temporary help, or to work overtime? Could an outside service be used? Would there be an increased error rate, which would increase error research and correction expenses? Might there be increased fraud loss, and fraud investigation costs? The cost to catch up, after the outage ends, is often a major extra expense factor. Beginning with the shortest significant duration, several estimates should be made, including a worst-case scenario.
- ***Lateness penalties***. If the application entails any lateness penalties, contractual, regulatory, or legal, an estimate should be made of the amount of the penalties that would be triggered by an outage, and the outage duration that would trigger the penalties.
- ***Public perception***. Public perception, as a catchall term, refers to the indirect effects of a service interruption. Staff morale and customer attitudes may be

affected negatively by the impression that the organization's risks are not being effectively managed. The value of publicly traded stock may be adversely affected, and pending mergers or stock offerings may be derailed. The cost of borrowing may increase.

Note that this procedure does not involve asking end users to estimate the maximum *tolerable* service interruption duration. Asking for a tolerable duration is *not* a useful question. An attempt by end users to establish the value of *tolerable* is simply an opinion or subjective judgment call. The basic risk assessment will establish the expected annualized losses. Then, an evaluation of the cost-benefit or ROI of alternative risk mitigation measures, implemented to reduce the number and duration of service interruptions, will identify the optimum IT system configuration, including measures to control service interruptions.

62.6.3.3 IT Systems Interruption Mitigation. IT system service interruption mitigation is a special case because of the effect of lost-data losses. This is why.

The performance of IT system continuity plans is often characterized using two parameters, referred to as *Recovery Time Objective*[14] (RTO) and *Recovery Point Objective* (RPO). RTO is the maximum service interruption duration that will result from the action of threats on an IT system. The RTO of the continuity plan in place for an IT system will determine the expected losses of an IT system that result from the impact of the service interruption threats. Decreasing the RTO will decrease the expected loss, but will increase the cost to implement and maintain the IT system continuity plan.

The RPO is a measure of the frequency with which an IT system backs up stored data. The more frequently data are backed up, the lower will be the lost data losses, but the greater will be the cost to install and maintain the backup provisions.

The risk-management task is to pick the IT system continuity plan (RTO-RPO combination) that has the most favorable cost-benefit. The *cost* is the cost to install and maintain a given RTO-RPO combination. The *benefit* is the reduction in expected loss the RTO-RPO combination will achieve, compared with the RTO-RPO combination currently in use. This is why the question: "What is your maximum 'tolerable' outage duration, that is, your RTO?" is worthless. A rational selection depends on both cost and benefit, not the subjective appeal of any particular RTO.

Typically, there is an interaction between RTO and RPO. For example, tape is probably less expensive as a backup storage medium than disk, so we would expect a tape-based IT system continuity plan to be cheaper than a disk-based plan. However, tape will restrict the range of RPOs that can be implemented, and may control the best achievable RTO, both of which will increase expected loss. There is another important point.

There is no inherent connection between the service interruption loss potential of an IT system and the cost of RTO-RPO combinations for the system.

For example, a small inexpensive system (one with low RTO and RPO costs) might be running a very time-critical, high-value function, such as currency trading, while a large system with high RTO and RPO costs might be analyzing the works of William Shakespeare for the local college, a low criticality task.

62.6.4 Risk Event Parameters. As noted in the prior section, during the determination of loss potential, a list of significant risk events can be constructed. The list should be examined critically to determine if additional event types should be

added to the list, bearing in mind the aggregation considerations discussed above and the defined scope of the project.

Once the risk events list has been completed, two parameters are estimated for each event as follows:

1. ***Occurrence rate***. There are several ways to estimate occurrence rates. If the event has had a relatively high occurrence rate in the past, typically more than once in 10 years, the organization is likely to have records of those events, from which the occurrence rate can be inferred with high confidence. However, there are two cautions. First, consideration must be given to any changes in the environment that would affect the occurrence rate. For example, the root cause of some system crashes may have been corrected, so that one would expect the future rate to be lower than in the past. Second, it is important to avoid double counting. For example, a system outage log may lump electric power failures into a catchall category: hardware failure. If electric power failure is included in the risk assessment as a separate risk event, its occurrences may be inadvertently counted twice.

   If the risk event is external to the IT facility, for example, river flooding and earthquake, one can usually find external sources of information about past occurrences. For example, one could use the NOAA National Hurricane Center's list of the United States' most intense hurricanes (www.nhc.noaa.gov/pastint.html) to determine if a given U.S. location has ever been impacted by a high-intensity hurricane. There are numerous Websites that can support a risk assessment. One also may find that past copies of local publications can help to determine the past local history. In some cases, there will be no meaningful past history for a risk event, particularly if the event is of recent origin, for example, e-business fraud and terrorism. In these cases, it may be possible to reason by analogy with other similar risk events to develop a reasonable estimate.
2. ***Outage duration***. The range of outage durations for each threat must be estimated in order to be able to estimate service interruption losses. If the outage duration is not the same for all the IT systems that the risk event impacts, one may either use an average value or define a separate risk event for each of the systems, keeping in mind the aggregation considerations discussed above.

Remember that until an initial overall risk assessment, like Exhibit 62.6, has been generated, the analysts do not know which of the input data are important. Consequently, it is counterproductive to devote effort to refining initial data. The wide range of values encountered in a typical risk assessment makes it clear that rough initial estimates of the risk parameters will be sufficiently accurate to identify the critical parameters.

62.6.5 Threat Effect Factors, ALE, and SOL Estimates. From the basic risk-assessment algorithms, it is possible to determine the risk event/loss potential pairs for which the Threat Effect Factor is greater than zero (see Section 62.6.2), and having done so, to estimate the value of the factor. It is then possible to use the risk data discussed above and the algorithms to calculate an estimate of the ALE and SOL for each of the pairs, and then to tabulate the individual ALE and SOL estimates to obtain a complete risk assessment for the IT organization. This is useful as a baseline because it provides an assessment against which to evaluate potential risk mitigation measures.

In other words, the benefit of a risk-management measure is the amount by which it reduces this baseline or current expected loss estimate.

In order to validate the assessment, it is necessary to identify the threat event/loss potential pairs that account for the majority of the expected losses. Typically, about 20 percent of the pairs cause 80 percent of the total expected loss. Reviewing the details of each input item, these questions should be asked: Were source data transcribed correctly? Are the assumptions reasonable? Were calculations made correctly? Does any threat event/loss potential pair account for so much of the expected losses that it should be disaggregated from other data items? Once the validation process is complete and indicated changes have been made, it is feasible to consider risk mitigation, but first one must gain acceptance of these data.

Field experience has disclosed a very powerful technique for gaining acceptance of estimates used in the risk assessment. First we gather data from knowledgeable people in the organization: line-of-business managers, IT specialists, audit, security, and so on. We use these data to make a preliminary baseline estimate of risks as described in the previous sections. We then convene a round table meeting of all the people who provided the data. We show them the results and the input data. This results in a useful discussion of the data, and tends to identify inputs based on faulty assumptions, personal biases, and misinformation. It is rather like an instant peer review. The agreed-upon changes to the input data are then used to make a final estimate of risks, which can be used to evaluate proposed mitigation measures.

62.6.6 Sensitivity Testing. The preceding discussion implies that a single value is to be estimated for each of the risk parameters included in the risk assessment. However, it may be thought that a parameter should not be limited to a single value, and that a Monte Carlo technique should be used to generate a set of ALE and SOL estimates. In a Monte Carlo approach, one generates random inputs into a model following specified probability distribution functions for each of the variables being studies. The outputs of the model then generate frequency distributions that allow one to evaluate the sensitivity of critical outputs to the precision of specific inputs. There are two obstacles to using this approach. Apart from the big increase in the complexity of the calculations, there is still a requirement to select a *single* risk mitigation strategy based on the set of ALE and SOL estimates. *It is not feasible to install a range of performance values for given mitigation measures* to match a range of loss estimates. As described previously, cost-beneficial risk mitigation measures are selected by calculating the baseline ALE, and then recalculating the ALE, assuming that the proposed mitigation measure has been installed. The difference between the two ALE estimates is taken to be a reasonable estimate of the return in the ROI calculation. Assuming two sets of 1,000 ALE estimates (without and with a proposed mitigation measure), which were generated by a Monte Carlo simulation, calculating all possible differences produces a set of 1,000,000 returns. Which one should be selected to calculate ROI?

Using the most likely value for each risk parameter, and generating only two ALE estimates, the difference between the two may be used to calculate ROI. Because the ROI typically flows from a relatively large number of parameters, the overall result will tend to average out any individual departures from the average values used, and so will be a fair representation of the actual risk losses that will occur.

However, if confidence in the accuracy of an estimate is extremely low, estimates using low, median, and high values of the parameters should be used to evaluate the effect on the overall results. Whenever a risk assessment includes a material low-confidence parameter, full disclosure should be made to the managers who review and

act on the results. Risk analysts should bear in mind that almost all other decisions about resource allocations, which senior managers are regularly called on to make, are based on uncertain estimates of future events. There is nothing unique about the uncertainties surrounding risk mitigation decisions.

62.6.7 Selecting Risk Mitigation Measures. The objective of the risk mitigation process is to identify the optimum set of mitigation measures. The first step is to address the intolerable SOL exposures as discussed in Section 62.5.3. The risk assessment and the evaluation of potential risk mitigation measures provide the raw material on which to base an implementation strategy for the risk exposures in the midrange. Considerations include:

The mitigation measures with negative ROIs can be discarded.

The remaining mitigation measures can be tabulated in three ways:

1. In descending order of net benefit: ALE reduction minus the implementation cost
2. In ascending order of implementation cost
3. In descending order of ROI

Based on other considerations, senior management selects mitigation measures from one of the three lists. Other considerations include the following:

- The availability of resources for risk mitigation may be limited, in which case some mitigation actions must be deferred until the next budgeting period. In extreme cases, the lowest cost measures may be selected.
- The mitigation of risk exposures that have a particularly undesirable effect on marketing considerations, as compared with other loss categories, may be given priority.

62.7 SUMMARY. This chapter has shown the tremendous advantage of having a detailed, quantitative assessment of the risk exposures leading to a quantitative evaluation of prospective mitigation measures. Senior managers will appreciate the advantage of using rational business judgment in place of seat-of-the-pants guesswork.

62.8 FURTHER READING

Anderson, E. *Business Risk Management: Models and Analysis*. Wiley, 2014.

Bhatia, M. *Handbook of Applied Risk Measurement*. Wiley, 2014.

Blundell, D. M., T. Cuthbertson, and J. Daniel. *Operational Risk Management*. Global Professional Publishing, 2013.

Coleman, T. S. *A Practical Guide to Risk Management*. Research Foundation of CFA Institute, 2011.

Crouhy, M., D. Galai, and R. Mark. *The Essentials of Risk Management*, 2nd ed. McGraw-Hill, 2013.

Gordon, L. A. *Managerial Accounting—Concepts and Empirical Evidence*, 6th ed. McGraw-Hill, 2005.

Lam, J. *Enterprise Risk Management: From Incentives to Controls*, 2nd ed. Wiley, 2014.

Louisot, J.-P. *ERM—Enterprise Risk Management: Issues and Cases*. Wiley, 2013.

McGill, W. L. *Math for Security*. CRC Press, 2014.

Moeller, R. R. *Executive's Guide to COSO Internal Controls: Understanding and Implementing the New Framework*. Wiley, 2013.

Mulcahy, R. *Risk Management Tricks of the Trade for Project Managers + PMI-RMP Exam Prep Guide*, 2nd ed. RMC Publications, 2010.

Siepman, F. *Managing Risk and Security in Outsourcing IT Services: Onshore, Offshore and the Cloud*. Auerbach, 2013.

Sweet, K. *Introduction to Security and Risk Management*. CRC Press, 2014.

Sunstein, C. R. *Risk and Reason: Safety, Law, and the Environment*. Cambridge University Press, 2002.

62.9 NOTES

1. The term *mitigation measure* is used here broadly to refer to all actions taken to manage risks, and may include logical controls over access to data, physical controls over access to IT hardware, facilities to monitor operations, selection and supervision of IT personnel, policy and procedures to guide operations, and so forth— in short, all of the security techniques described in this *Computer Security Handbook*.
2. Donn B. Parker, "What's Wrong with Information Security—and How to Fix It." Lecture at Naval Postgraduate School, April 28, 2005, www.youtube.com/watch?v=RW9hOBCSy0g
3. Another way of expressing this concept is as follows: Perfect security is infinitely expensive, and so it is not a rational goal of a risk-management program.
4. As a practical matter, it may become obvious that in some situations the differences between the members of a group of like IT systems are too small to be significant, and that a single set of security measures can be adopted for all of the systems.
5. "Expected loss" is a shorthand term for the losses an organization can reasonably expect to experience given its risk environment and the potential for loss of its functions and assets. It is the sum of the individual ALEs.
6. This author has made it a practice at the end of each interview to ask: "Is there something I should have asked about?" or words to that effect. In most cases, new information emerges as a result.
7. The plot was generated automatically by CORA® (Cost-Of-Risk Analysis), a risk-management software system.
8. D. McNeill, "With Fukushima Nuclear Plant Still Leaking, Japan Clean-Up Bill Soars to $50bn." *The Independent*, July 24, 2013, www.independent.co.uk/news/world/asia/with-fukushima-nuclear-plant-still-leaking-japan-cleanup-bill-soars-to-50bn-8730832.html
9. Return on investment is the ratio of the return received from an investment to the investment itself. If a bank pays back $1.05 one year after $1.00 has been deposited in the bank, the ROI is roughly 5%. To determine ROI accurately, one must take the ratio of the present value of the return to the present value of the investment.
10. A given measure may address more than one threat, so the total of the individual reductions in expected loss are used to evaluate the ROI of the measure.
11. This is an example of the advantage of aggregation. Rather than estimating the loss potential for individual physical item, it is much more efficient to aggregate the items into large classes. It may also be feasible to estimate loss potential on the basis of replacement cost per unit area.

12. The term *liability* is used here as a shorthand for *exposure to liability*. For example, an IT service performed for others might expose the organization to a liability action if the service was performed incorrectly resulting in a client loss. This potential risk exposure is referred to here simply as a liability.
13. The term *denial of service* implies that service interruptions are the result of deliberate attacks on an IT system, but, of course, there are many other threats that cause service interruptions. For this reason, the term may cause misunderstanding, and so should not be used.
14. The use of the word *objective* is unfortunate in this context. It would be more precise to use a term like *provision of performance* to make clear that the quantity is what a contingency plan is expected to achieve . . . not a vague goal. However, the terms *RTO* and *RPO* are widely used, so we will use them here.

CHAPTER 63

MANAGEMENT RESPONSIBILITIES AND LIABILITIES

Carl Hallberg, M. E. Kabay, Bridgitt Robertson, and Arthur E. Hutt

63.1 INTRODUCTION 63·1
- 63.1.1 Role of Management 63·2
- 63.1.2 CISO 63·2
- 63.1.3 Information Security Integrating into Strategic Vision 63·4
- 63.1.4 Net Present Value of Information Security 63·5
- 63.1.5 Case Study: Veterans Affairs 63·5

63.2 RESPONSIBILITIES 63·10
- 63.2.1 Policy Management 63·12
- 63.2.2 Motivation 63·12
- 63.2.3 Supervision 63·14
- 63.2.4 Judgment and Adaptation 63·15
- 63.2.5 Management Failures 63·16
- 63.2.6 Risk Management 63·18

63.3 LIABILITIES 63·19
- 63.3.1 Stakeholders 63·19
- 63.3.2 Due Diligence of Care 63·20
- 63.3.3 Downstream Liability 63·20
- 63.3.4 Audits 63·22

63.4 COMPUTER MANAGEMENT FUNCTIONS 63·23
- 63.4.1 Planning for Computer Security 63·23
- 63.4.2 Organizing 63·23
- 63.4.3 Integrating 63·24
- 63.4.4 Controlling 63·25

63.5 SECURITY ADMINISTRATION 63·25
- 63.5.1 Staffing the Security Function 63·25
- 63.5.2 Authority and Responsibility 63·26
- 63.5.3 Professional Accreditation and Education 63·27

63.6 CONCLUDING REMARKS 63·28

63.7 FURTHER READING 63·29

63.8 NOTES 63·29

63.1 INTRODUCTION. This chapter reviews the critical roles of management in establishing, implementing, and maintaining information security (INFOSEC) policies in the modern enterprise. It also reviews some of the risks to management personnel in failing to ensure adequate standards of INFOSEC.[1]

63.1.1 Role of Management. Organizations are unequally affected by the risk of loss. In certain government computer installations, matters of national security are at stake, and the measures required to protect such facilities are elaborate and costly. At the other end of the spectrum are computers used exclusively for word processing of unclassified materials, which are not connected to networks; these require few security precautions except for file backup and antivirus software. This chapter and this *Handbook* do not address either of these extremes, but rather the bulk of user organizations in business, government, and universities, where concentrations of information assets and dependence on computers create an exposure to loss, and managers must balance security with cost effectiveness and common sense while convincing colleagues with a superficial knowledge of security to pay attention to their recommendations.

Management provides the essential framework for accomplishing technical work. Whether it is drawing up a security policy, enforcing such policies, training the individuals who will implement and enforce those policies, or proposing a budget to get all this done, management plays an essential role. Information technology (IT) managers ensure the consistent functioning of the organizational computing environment. Ideally, they also provide insights and guidance to upper management in strategic planning to take advantage of new opportunities.

Many organizations, regardless of size and history, have heterogeneous networks with many different operating systems running different applications and serving many purposes and clients. Web servers sit next to email servers, which connect to outside networks, which then rely on connections to more than one third-party corporate network. The rapid pace of technological change impedes the IT manager's ability to keep the enterprise IT infrastructure running smoothly.

A key function of IT managers is loss avoidance by managing risk intelligently, in order to reduce the likelihood of trouble in the IT sector and to reduce the costs of coping with such trouble.[2] In a broader sense, this entire *Handbook* is designed to support management in these efforts.

IT managers must focus constantly on enabling business functions. Given that most enterprises do not consider security as their main task, IT managers must ensure that INFOSEC policies and technology support, rather than hinder, the principal business of the enterprise. It is equally important to impress this philosophy of service to the strategic and operational goals of the enterprise on all members of the security staff. Many of these members may have developed their careers entirely in the technical sectors, and they may not have developed a service orientation toward their nontechnical fellow workers, and toward the enterprise itself.

63.1.2 CISO. In recent years, INFOSEC has been growing in visibility as a major business concern. IT managers are now being joined by INFOSEC managers. For example, a major bank responded to infiltration of its systems in the 1990s by naming a chief INFOSEC officer (CISO) reporting at the same level as officers such as the chief information officer (CIO), chief financial officer (CFO), chief operations officer (COO), chief technology officer (CTO), and chief executive officer (CEO). Together, the CIO and the CISO face complex responsibilities, and are increasingly visible in the corporate infrastructure.

Today, CISOs are increasingly found in all types of corporations. They may report to the CIO, or they may have a position on an equal level with the CIO. Some of the functions in a security organization are suggested in the hierarchy that follows. (These

functions do not imply one team or one person per function, but simply illustrate the range of functions a CISO must coordinate.)

- CISO
  - Security policy team
    - Development
    - Awareness
    - Education
    - Update management
  - Access controls
    - Data classification functions
    - Identification, authentication, and authorization functions
  - Intrusion detection and prevention
    - Outward-facing intrusion detection systems
    - Gateway security devices
    - Malware control team
  - Security assessment
    - Security auditing
    - Penetration team
  - Engineering coordination
    - Software quality assurance coordination
    - Patch management review
    - Network security oversight
    - Control systems and logging management
  - Security incident response team
    - Planning and rehearsals management
    - First line response
    - Data gathering and forensic analysis
    - Coordination with internal functions (legal, public relations, human resources, etc.)
    - Coordination with law enforcement

One of the factors in this increased visibility of the CISO is the rapid growth in the popular press of technology issues coverage and of INFOSEC breaches in particular. For example, distributed denial-of-service attacks, malware infestations, breakdowns of privacy policies, theft of credit card numbers, loss of control of large volumes of personally identifiable information, identity theft, Website defacements, spam, phishing, and the other forms of abuse to information systems, described in several chapters of this *Handbook,* have kept the public aware of INFOSEC and its breaches.[3]

In some sense, publicity has helped the IT world, because some IT managers now have increased ammunition with which to argue for management support of the security function. Publicity has increased visibility, however, so that even a minor breach of security may spark an overreaction in the upper echelons of the enterprise. Indeed, in

California, it is now law (SB 1386) that companies that suffer from some forms of INFOSEC breaches must quickly advise their customers of these events.[4]

63.1.3 Information Security Integrating into Strategic Vision. IT security managers are now able to help their businesses achieve strategic goals and to further their responsibilities by showing that security involves more than just protecting assets, it is also potentially a business enabler. In the absence of a secure IT infrastructure, businesses may be slow to enter a given market for fear of losses and liabilities resulting from security breaches. However, with security integrated into the corporate culture, these enterprises can confidently enter new arenas and convince their potential clients that their data will be safe. The authors have personal experience from consulting assignments in which clients stated that their own potential customers were reluctant to do business with them because their security measures were inadequate.

George Lin describes the benefits of a business-focused IT organization:

> In a business-focused IT organization … IT staff are more interested in understanding the business—and the people, processes, and organization that make up the business—than they are in the technology for its own sake. These staff members possess a balanced set of soft skills, and business acumen as well as analytical and technical skills. They think about business processes first and technology second. The solutions they propose tend to be more complete, often involving people, process, and organization as well as technology, which is seen as simply a tool. Interacting with such an organization yields an experience similar to that expected from Big Four consulting firms, but with the added intimacy and insights only an internal organization can provide.
>
> More importantly, IT/business alignment becomes a nonissue because it is natural to a business-focused IT organization, the very DNA of its staff. There is no need to expend effort specifically on IT/business alignment because the business-focused IT organization thinks and acts like the business that it serves, and is an integral part of that business.[5]

In this framework, INFOSEC becomes not an afterthought, preventing loss, but part of the design strategy, which will help realize gain. This is a much more powerful, and positive, view of INFOSEC's role in the business world than is the traditional one. Using INFOSEC principles in the design of a product or service generally helps to create a more robust product or service, one that will be less prone to risk.

The National Institute of Standards and Technology has articulated a roadmap for integrating security considerations into the system development life cycle (SDLC)[6] (with addition of three more security objectives in compliance with the Parkerian hexad described in Chapter 3 in this *Handbook*):

> A general SDLC includes the following five phases: *initiation, acquisition/development, implementation/assessment, operations/maintenance, and sunset (disposition)*. Each of these five phases includes a minimum set of security tasks needed to effectively incorporate security in the system development process. Including security early in the information SDLC will usually result in less expensive and more effective security than adding it to an operational system.
>
> The following questions should be addressed in determining the security controls that will be required for a system:
>
> - How critical is the system in meeting the organization's mission?
> - What are the security objectives required by the system, e.g., integrity, confidentiality, and availability [and control, authenticity, and utility]?
> - What regulations and policies are applicable in determining what is to be protected?
> - What are the threats that are applicable in the environment where the system will be operational?

Integrating a security methodology into a product or service helps to defines roles and responsibilities and helps to manage expectations.

63.1.4 Net Present Value of Information Security. A paper from the late 1990s articulated the positive value of INFOSEC in a competitive economy. The summary included these conclusions:

> The news is that we can now go beyond gibbering and fist shaking. In today's e-commerce environment, effective information security can actually increase business and increase profits, not merely reduce risk. …
>
> Tom Nelson, VP and Chief Strategy Officer of AtomicTangerine, has defined the Net Present Value of Information Security (NPVSec) as follows: "NPVSec is the value protection and value creation that is realized when barriers to e-Business are removed through mechanisms that ensure business integrity, service availability, and customer/consumer confidentiality and privacy. Value creation examples include: new distribution channels, new revenue streams, new business models, among others."
>
> In other words, instead of viewing information security solely as a risk-avoidance measure—like a kind of insurance policy that never actually pays anything back—we are forced by the nature of e-business to accept that security actually *supports* and *enables* e-business.
>
> As we have seen … e-business has brought security to the forefront of strategic thinking for successful businesses. Business leaders can no longer tolerate the view that security is an add-on feature relegated to the end of the design process. Security is a process, not a product or a state; security affects every e-businesses bottom line in a positive way. Security is no longer a cost center; it's part of your repertoire for meeting the legitimate needs of your public. Instead of seeing security as solely the purview of the technical staff in your organization, you should ensure that your marketing and public relations departments are well versed in the principles of information security and can communicate effectively to an anxious public about the measures you are taking to safeguard your customers' privacy and their money.[7]

Another perspective on the value of integrating information security into a corporate culture argues that "the responsibility for protecting the organization's information assets is no longer restricted to the CSO [chief security officer]. Every employee has a responsibility to help protect the proprietary data they are entrusted with."[8] The authors write that:

> This approach:
>
> - Improves individual leadership through collaboration and team building
> - Strengthens corporate allegiance
> - Helps to provide a sense of community
> - Builds an intelligent workforce where every employee is recognized and watched by their peers for their individual contribution to the team's strength. …
>
> In addition, embedding security deep into every process of the corporation, starting at the executive level, is an extremely effective method of taking one problem and leveraging the success of that solution to correct other problems.

63.1.5 Case Study: Veterans Affairs. The next summary of a specific case of management failures will serve to help readers think about the management responsibilities for information assurance.[9]

63.1.5.1 Announcement without Taking Responsibility. In March 2007, *Network World* writer Jon Brodkin wrote an excellent analysis of 10 letters

informing victims of data theft or loss of control of personally identifiable information (PII) that their data might be compromised.[10] He pointed out that almost all of the letters failed to express any responsibility for the loss of control over data stored on unencrypted disks that were lost or stolen or for poorly secured Websites that posted PII without protection, or with poor protection. Perhaps staff attorneys warned the public relations officials to avoid any implication of responsibility to avoid contributing anything that would exacerbate their liability in potential lawsuits. Passive voice is often used to shift responsibility from specific agents to the great gaseous cloud of the unnamable and unblamable. The classic example is "Mistakes were made."

In 2007, this letter was sent to physicians affected by a security breach.

DEPARTMENT OF VETERANS AFFAIRS

1615 Woodward St.

Austin, TX 78772

—–, MD

Dear —–, MD:

I am writing to you, as the Director of the Veterans Integrated Service Network (VISN) 7 in Atlanta, Georgia, to inform you that I have been notified that a portable computer hard drive used by an employee of the Birmingham Veterans Affairs (VA) Medical Center is missing. This portable hard drive was used to back-up information contained on a VA employee's office computer, related to research projects with which the employee was involved. A file on the portable hard drive included information from the Unique Physician Identification Number (UPIN) Directory dated 2004, which includes demographic information and identifiers, such as the UPIN, dates of birth, state license numbers, business addresses, and employer identification numbers (EIN). In the case of your information, we believe the EIN was your Social Security Number. This file was obtained by VA from the Centers for Medicare & Medicaid Services (CMS) for the purpose of conducting research on veterans' health care.

The Birmingham VA Medical Center has conducted extensive physical searches and has involved local police and Federal investigative resources, and a reward is being offered; however, the hard drive remains missing. To prevent further security breaches or losses, we have taken immediate measures to protect the integrity and security of all personally identifiable information including prohibition of the use of external drives and the required encryption of personally identifiable information when authorized distribution is required.

An independent risk analysis was conducted as required by law, and risk mitigation recommendations are being implemented immediately. VA will contact you shortly by mail to offer a credit monitoring service at no cost to you. In the mean time, one precaution we recommend is for you to request a free credit report from one or more of the three national credit bureaus by calling the toll free number 1-877-322-8228. The credit bureaus may also be contacted at:

Equifax
P.O. Box 740241
Atlanta, GA 30374
1-800-685-1111
Experian
P.O. Box 9554
Allen, TX 75013
1-888-397-3742

TransUnion
P.O. Box 2000
Chester, PA 19022
1-800-916-8800

> More information about credit protection, including placing a "fraud alert" on your accounts, is available by calling the Federal Trade Commission at its toll free number, 1-877-438-4338, or by visiting its Website, http://www.ftc.gov/
>
> If you have questions concerning this letter, the Birmingham VA Medical Center has established a dedicated call center to answer your questions. Please contact us toll free at 1-877-xxx-xxxx from 6:00 am to 9:00 pm CT, or email us at < address suppressed >.
>
> We at VA take information security and privacy very seriously. We apologize for any inconvenience or concern this situation may cause, but we believe it is important for you to be fully informed of any potential risk to you.
>
> Sincerely,
>
> [digitized signature]
>
> Lawrence A. Biro
>
> Network Director, VISN 7

63.1.5.2 Initial Problems. On May 3, 2006, a career civil servant at the Department of Veterans Affairs (VA) violated official policy by taking computer disks containing PII about 26.5 million veterans home with him. The disks were stolen from his home.[11] Two weeks after officials learned of the theft, the VA disclosed the incident to the public and set up a Website and an 800-number to provide veterans and with information and a channel for reporting possible identity theft.[12]

The USA.gov Website put up a page called "Latest Information on Veterans Affairs Data Security"[13] with answers to frequently asked questions; the VA itself also continued issuing press releases.[14]

In early June 2006, the VA announced that the stolen data might include PII about up to 1.1 million active-duty troops, 430,000 members of the National Guard, and 645,000 members of the reserves.[15] Reactions from a coalition of veterans groups were immediate: They launched a class-action lawsuit demanding full disclosure of exactly who was affected by the theft and seeking $1,000 in damages for each victim.[16]

The VA struggled to cope with the bad publicity and potential legal liability resulting from the May theft. On May 26, 2006, Secretary of VAR James Nicholson issued a Directive to all VA supervisors, in which he wrote:

> Having access to such sensitive information brings with it a grave responsibility. It requires that we protect Federal property and information, and that it shall not be used for other than authorized activities and only in authorized locations. As managers, supervisors, and team leaders it is your responsibility to ensure that your staff is aware of and adheres to all Federal and VA policies and guidelines governing privacy protected material. I also expect each and every one of you to know what sensitive and confidential data your subordinates, including contractors, have access to and how, when and where that data is used, especially in those cases where it is used or accessed off-site.[17]

On May 30, 2006, the VA fired the analyst "responsible for data loss" and announced changes in the administration of INFOSEC in the organization.[18] The press release made no mention of who was responsible for allowing anybody to store unencrypted PII on VA computers or media.

Coincidentally, at the end of May, the Government Accountability Office (GAO) issued a report: "GAO-06-612: Homeland Security: Guidance and Standards Are Needed for Measuring the Effectiveness of Agencies' Facility Protection Efforts."[19] The

report specifically named the VA as requiring "guidance and standards for measuring performance in federal government facility protection."

On June 21, 2006, the VA announced that it would provide free credit monitoring for everyone affected by the data theft in May.[20]

But worse was yet to come.

63.1.5.3 Systematic Management Failures. On June 14, 2006, Linda D. Koontz, Director, Information Management Issues, and Gregory C. Wilshusen, Director, Information Security Issues of the Government Accountability Office of the United States, offered testimony before the Committee on Veterans' Affairs, House of Representatives. The GAO report on their analysis and recommendations later appeared as GAO-06-866.[21] Highlights of their analysis included these comments:

> For many years, significant concerns have been raised about VA's information security—particularly its lack of a robust information security program, which is vital to avoiding the compromise of government information, including sensitive personal information. Both GAO and the department's inspector general have reported recurring weaknesses in such areas as access controls, physical security, and segregation of incompatible duties. The department has taken steps to address these weaknesses, but these have not been sufficient to establish a comprehensive information security program. For example, it is still developing plans to complete a security incident response program to monitor suspicious activity and cyber alerts, events, and incidents. Without an established and implemented security program, the department will continue to have major challenges in protecting its information and information systems from security breaches such as the one it recently experienced.

Two related reports appeared about a week later with specific comments about the May 2006 data breach[22] and about the overall challenges facing the VA and the Department of Defense (DoD) in protecting PII of active-duty and retired military personnel.[23]

At the end of June 2006, the laptop and external hard drive stolen on May 3 from the consultant's home were recovered. Forensic examination suggested that the data had not been accessed. This good news suggested that the disaster might blow over.

It was not to be.

The Inspector General (IG) of the VA, George Opfer, released a report on July 11 severely criticizing senior managers of the VA for their lackadaisical response to the original theft of unencrypted PII. The inadequate data security policies had not yet been corrected.[24] VA secretary James Nicholson responded to the IG's report with assurances that the agency had "embarked on a course of action to wholly improve its cyber and information security programs."[25]

63.1.5.4 Continued Problems. On Monday, August 7, 2006, Secretary Nicholson announced that a Unisys subcontractor working for the VA offices in Philadelphia and Pittsburgh had reported that his desktop computer was missing. The computer contained PII for 18,000 and possibly up to 38,000 veterans.[26]

A week later (August 14), the VA announced that it would spend $3.7 million on encryption software and would encrypt data on all the department's computers and external data storage media or devices. Installation would begin Friday, August 18.[27]

In mid-September, the stolen Unisys desktop computer with VA data was located and a temporary employee working on subcontract to Unisys was arrested and charged in the theft.[28]

In October 2006, the Congressional Committee on Oversight and Government Reform published a report on data losses in U.S. Government agencies since January 1, 2003.[29] There were 788 incidents in 19 agencies—in addition to hundreds of incidents at the VA. The report's findings included these bald assertions:

1. Data loss is a government-wide occurrence. …
2. Agencies do not always know what has been lost. The letters received by the Committee demonstrate that, in many cases, agencies do not know what information has been lost or how many individuals could be impacted by a particular data loss. Similarly, agencies do not appear to be tracking all possible losses of personal information, making it likely that their reports to the committee are incomplete. For example, the Department of Justice reports that, prior to the May 2006 Veterans Administration data breach, "the Department did not track the content of lost, stolen, or otherwise compromised devices."
3. Physical security of data is essential. Only a small number of the data breaches reported to the Committee were caused by hackers breaking into computer systems online. The vast majority of data losses arose from physical thefts of portable computers, drives, and disks, or unauthorized use of data by employees.
4. Contractors are responsible for many of the reported breaches. Federal agencies rely heavily on private sector contractors for information technology management services. Thus, many of the reported data breaches were the responsibility of contractors.

Alas, the best-laid plans of VA administrators gang aft agley, and on October 31, 2006, VA officials informed 1,400 veterans that their PII had been on unencrypted data disks sent by mail from the VA clinic in Muskogee, Oklahoma, on May 10, June 10, and July 10, which were lost. A spokesperson for the hospital explained the three-month delay as being due to the "wait for officials in Washington to approve the wording of the letter." Approval arrived October 26. There was no explanation of why the data were unencrypted, nor why two additional disks were mailed out after the May 10 disk was lost. A report on this incident dated November 3, 2006, by Rick Maze in the *Federal Times*[30] also indicated that a laptop computer from the VA hospital in Manhattan was stolen on September 8 from a computer locked to a cart in a locked room in a locked corridor—and that the data on the stolen machine was deliberately not encrypted despite policy because "a decision had been made not to encrypt data being used for medical purposes."

And more was to come in February 2007.

63.1.5.5 Analyses and Responses. On Friday, February 2, 2007, Secretary of Veterans Affairs Jim Nicholson announced that a VA employee in the VA medical center in Birmingham, Alabama, had reported an external hard drive as missing on January 22. According to Representative Spencer Bachus (R-AL), the backup hard drive contained PII on up to 48,000 veterans—and despite VA regulations promulgated in 2006, as many as 20,000 of those records were not encrypted.[31] A week later, the VA admitted that the hard drive actually contained PII about 535,000 patients and 1.3 million doctors.[32] It was that loss that led to the letter quoted in the first part of this section.[33]

A few weeks later, the GAO released the closest thing to an exasperated blast of which government workers are capable. In testimony before the Subcommittee on Oversight and Investigations, Committee on Veterans' Affairs, House of Representatives on February 28, 2007, GAO Director of Information Security Issues Gregory C. Wilshusen presented a report entitled "Veterans Affairs Needs to Address

Long-Standing Weaknesses."[34] The summary on page 2 of the PDF file includes this commentary:

> For many years, GAO has raised significant concerns about VA's information security—particularly its lack of a comprehensive information security program, which is vital to safeguarding government information. The figure below details information security weaknesses that GAO identified from 1998 to 2005. As shown, VA had not consistently implemented appropriate controls for (1) limiting, preventing, and detecting electronic access to sensitive computerized information; (2) restricting physical access to computer and network equipment to authorized individuals; (3) segregating incompatible duties among separate groups or individuals; (4) ensuring that changes to computer software were authorized and timely; or (5) providing continuity of computerized systems and operations. The department's IG has also reported recurring weaknesses throughout VA in such areas as access controls, physical security, and segregation of incompatible duties. In response, the department has taken actions to address these weaknesses, but these have not been sufficient to establish a comprehensive information security programs. As a result, sensitive information has remained vulnerable to inadvertent or deliberate misuse, loss, or improper disclosure. Without an established and implemented security program, the department will continue to have major challenges in protecting its systems and information from security breaches.

In early March 2007, the VA reacted to the January 22 loss of the portable hard drive. CIO Robert Howard promulgated a policy restricting the use of portable data storage devices. Only flash drives smaller than 2 GB—and only those issued by the VA's CIO office itself—would be permitted on the VA network or computers. Encryption would be used throughout the system, just like the assurance issued in August 2006 about spending $3.7 million on encryption tools.[35] In addition, the CIO announced sweeping changes in security administration, with promotion of five deputy CIOs to the rank of assistant secretaries for these functions: application development, information security, operations and maintenance, resource management, and strategic planning.

As of late May 2007, federal agencies announced that they would stop storing Social Security numbers and other PII wherever possible.[36]

For an extensive compilation of additional cases of interest, see the "Chronology of Data Breaches: Security Breaches 2005–Present," managed by the Privacy Rights Clearinghouse.[37]

63.2 RESPONSIBILITIES. Modern information processing philosophy transcends the boundaries of the computer room and demands the consistent delivery and collection of data to and from remote sections of the organization. Once this operating mode is extensively used, the company becomes dependent on its continued availability. Any shutdown or cutoff in service can be disastrous, unless the company is able promptly to revise operations, in order to continue vital activities. Central computers, local servers, workstations, and the network are essential to the operating environment.

Computer security includes the protection of confidentiality, control, integrity, authenticity, availability, and utility of information.[38] Computer security encompasses the total infrastructure for maintenance and delivery of information, including physical computer hardware, supporting equipment, communication systems, logical processes defined by software, and the human factors that support and possibly threaten this infrastructure.

Computer security is inseparable from the basic structure of the information processing system; one cannot and should not design a system without including security as an underlying strategy. The objectives and properties of secure systems must be considered collectively from their inception.

Software, systems, and networks should be designed to ensure information protection that corresponds to business needs. There are many steps that can be taken to do this, but managers must place a high priority on making it happen. Managers are in a position to see what the business needs actually are, much more so than the typical INFOSEC analyst or engineer. A prime example is e-commerce Website management. No reputable site allows transactions without minimal SSL encryption. A savvy INFOSEC manager will track all information from the setting up of the SSL tunnel, all the way through the shipping process, and beyond into the storage of information gathered during any transaction, and ensure that security is not lacking at any juncture in the process. If the data are not securely stored while sitting in a database, the whole process falls apart.

The ability of a system to ensure accuracy, reliability, and confidentiality is a basic building block. The collective system (hardware, software, communications, and people) must be able to maintain and process data correctly, and move traffic (transactions, inquiries, commands, etc.) from its origins to the intended destinations, without unauthorized modification or disclosure and without misrepresentation, forgery, or other breaches of authenticity. Reliable performance is essential. Any failures should be orderly and predictable, with adequate detection methods to provide timely evidence of failure and to permit prompt corrective actions. Hardware/software limitations should be known and documented, as should load limits. Bounds checking, at a minimum, should be implemented in all software to avoid buffer overflows. Any inputs that are not checked should be documented, and measures should be put in place to detect possible abuses.

Suspicious user activity should be detectable through appropriate analysis of logfiles. The network and its data should be protected from contamination or outside interference using appropriate gateway security devices and other intrusion prevention systems. Firewalls must be present, updated, and frequently reviewed, along with intrusion detection systems and intrusion prevention systems (IDS/IPS), and all logs should be centralized and scrutinized. Measures should be put in place to ensure that data transmitted are the same as the data received and that error correction is used wherever possible.

Service-level agreements should define goals for efficient response and adequate capacity, in order to support acceptable performance. Systems should be able to recover quickly from either short-term or long-term disruptions. Backup measures, including data backup and equipment backup[39] and tested incident response, contingency, and recovery plans, should be in place.[40] Appropriate cost recovery measures should also be in place.[41] Prevention of harm is of primary importance.

Single points of failures should be avoided. Avoid dependence on single equipment devices or single communications pathways. Web servers should be clustered, and databases should be replicated across various servers and placed in different geographic regions. Avoid overloading the network at peak activity periods. Provide environmental backup (redundant power, air-conditioning, heating and other support systems, limited equipment access, etc.) to reduce other exposures.[42]

Management should have the capability to limit who can access the systems, how much capacity can be used for each purpose or function, what purposes are allowed, what data are accessible and transmittable for each user, and what connections can be made. Technical measures include access control (system access, resource restrictions), logon password control, alternative identification methods (personal tokens, digital signatures, biometric authentication), callback connections, network isolation from public networks, firewalls, and physical security of system components.[43]

Administrative measures include publication of policies, standards, and guidelines; screening of personnel; security awareness training; system change control procedures, including security criteria in system design; and monitoring system activity and quality.[44]

In addition, management should use the audit function as an independent arbiter to measure compliance with policies, standards, and guidelines as well as to assess the adequacy of technical protection. The system and its network should preserve and display evidence of use, behavior, and content, and should record deviations from expected use.[45]

Managers are responsible for specific tasks or functions to the extent that they make decisions about business processes and suffer the consequences or reap the benefits of those decisions.

63.2.1 Policy Management. Managers spend more time on people issues than on technical issues. The right people must be hired for each position.[46] Employees who do not work responsibly and competently must either be brought up to the proper standard or let go. Employees who do their jobs properly must be kept satisfied, lest they move to another organization. Management in the IT world must ensure compliance with corporate policies.[47] Compliance with policies includes motivation, supervision, judgment, and adaptation.

63.2.2 Motivation. Employees need motivation to pay attention to INFOSEC policies, which often are perceived as a nuisance, interfering with the fundamental goals of the enterprise. Upper management, in particular, must set an example by following enterprise policies; when top managers are seen to ignore policies, the people reporting to them quickly imitate their behavior, and the problem spreads throughout the organization. For example, if the CEO refuses to wear a picture badge, the vice presidents (VPs) will quickly follow suit because they will associate *not* wearing badges with high status. Similarly, the directors reporting to VPs will start dropping their badges, and so within a few months the entire hierarchy will be convinced that no one but stock clerks should wear badges. Sometime later, the stock clerks will be resisting badges too.

According to *Computerworld* writer Mary Pratt, writing in 2006:

> The risk that employees pose is significant. They can fall prey to social engineering, a fancy term for being conned. They can ignore company policy by failing to encrypt sensitive data. Or they might install unauthorized software that can corrupt the system.
>
> Think you're well protected? Recent findings from the Computing Technology Industry Association might convince you otherwise. In this year's CompTIA information security study, 59% of the organizations surveyed indicated that their latest security breaches were the result of human error alone. That's up from 47% last year.[48]

A survey by NFO Prognostics highlighted that 66 percent of the companies surveyed believe that staff training or certification has improved their IT security through increased awareness and proactive risk identification.[49]

Since people can be the company's greatest strength against attacks, training employees should be a high priority in any organization. At Stop and Shop, in February 2007, employees noticed suspicious activity; they saw four people tampering with the keypads for credit/debit card authorization units. The employees notified police at once, and the suspects were arrested. That kind of quick thinking and commitment can make

an enormous difference in the effectiveness of security systems of all kinds—and it depends on management support and motivation.[50]

Some managers think that punishment is the only motivation that can change behavior; but everything known about human psychology shows that reward is more powerful.[51] Reward is not limited to salary increases and bonuses; sometimes the most effective way to keep IT employees satisfied and productive is to provide training. The IT world is evolving, and the demand for competent staff has never been higher. Most employees want to feel that their employers value their services and that they are worth an investment to improve staff competence. Training employees, and providing challenging opportunities for the exercise of intelligence, serves the interests of both employer and employee. See Section 63.5.3 for more details about education.

Challenging employees in other ways can be highly motivating. For example, managers can encourage staff to prepare and deliver presentations at internal meetings and at security conferences. Some examples of useful conferences are SecureWorld-Expo, InfoSecWorld, and NetSec. Motivated employees can lead special interest group discussions at conferences (such as RSA, SANS, etc.) and so develop a Web of relationships that promote sharing of knowledge, and that enhance their self-image. There are also security contests to test their skills, such as the Honeynet forensics contest.

The Honeynet Project provides archives of past security challenges that allow security staff to challenge their technical skills, as well as to learn and teach new tricks of the trade. The Honeynet Project posts these challenges on the Web,[52] and tasks would-be contestants to unravel what the hacks were and how they were performed. Along with finding out what had happened, judging was also based on how much information was uncovered and how this information was communicated. These challenges have been solved, but they are still valuable lessons for any INFOSEC staff. Many other challenges pop up frequently and easily can be found using Google.

Many journals and Web-based magazines are ready to accept articles written by professionals in the field. Not only does writing solidify employees' knowledge and build their own confidence, but it also instills confidence in the entire team. Writing and teaching help the enterprise, as well as its clients and partners, to view the security team as a real benefit, at the same time that the organization develops a strong reputation for excellence in security.

Another way for management to make the security team more cohesive is to build camaraderie. A monthly pizza party or an occasional outing to a sports event can do a lot to ease stress and to help everyone to know each other. Having good friends at work can reduce turnover and motivate employees to do their best, not only for the rather abstract goal of doing good for the enterprise, but also because of a commitment to their colleagues.

To help encourage a higher level of expertise and to establish a feeling of belonging, periodic "brown bag" lunch sessions can be effective. At each session, a different team member can be designated to present an informed talk or to lead a discussion on a topic of general interest. There could even be an informal call for papers, giving team members an opportunity to present valuable information in a professional manner.

At the same time, a little friendly competition can help as well. For example, if the enterprise has a training network that can be subdivided into several subnets, managers can organize a "tiger-team challenge." Each team will be responsible for securing a subnet or host and then given the opportunity to break into the subnet or host of another team. The winners would get both a reward and a responsibility. The reward can be as simple as pizza for the team or a modest trophy. The responsibility would be to present the exploits used and the ways to secure against them to the rest of the teams.

Another way to build the team is to support staff in choosing the areas in which they want to excel. There may be some overlap, but overlap can be good. Clearly, if there are two firewall experts, the enterprise is less vulnerable should one of them be absent. Furthermore, the firewall experts can provide better-quality assurance by discussing alternatives when planning a change and better-quality control by checking each other's work. It can also help to have the security team members switch roles periodically, both to ensure that no one gets into a rut and to contribute to the challenges presented when a change is made.

Last, almost nothing is more infuriating than being expected to accomplish a task without the necessary resources. Adequate time to get a task done is always an important issue, as many people have several tasks to do, and each of the tasks may be seen as critical by someone. A manager should be willing to give the team all of the resources it needs to complete a task in the allotted time, or it should be made clear to other departments, or to higher management, what the realistic expectations are for the completion of such tasks, if required resources are lacking. If expectations are properly managed, there should be less conflict and fewer problems.

It is useful to view management not as separate from the IT security team, but as an integral part of it. This allows management to contribute directly toward the employees' enthusiasm as well as to detect early warning signals of impending trouble. A manager who does not spend time listening to the team members, or who does not understand what their jobs involve at a technical level, will not be respected by the team. Lack of respect will block communication and keep the department from becoming a solid, effectively functioning unit.

63.2.3 Supervision. What we know about damage to computer systems indicates that errors and omissions are a major source of harm. Poorly trained employees make mistakes, but so do trained employees who have become careless. Managers should examine performance records as a normal part of their supervision of the security team. In particular, every incident that damages production systems should be analyzed to identify the *reasons* for the event. Careful technical support records and log files can help the team spot the crucial weaknesses, whether technical or human, that allowed compromise of the damaged systems.[53]

Analysis of security breaches of all kinds may reveal that certain employees are associated with unusually high or unusually low frequencies of particular problems. Careful analysis of both types of extremes can be helpful in spotting weaknesses for remediation and strengths from which to learn, so that the knowledge can be spread across the entire unit. However, managers must not assume that disproportionate numbers of problems are necessarily caused by the employees involved; for example, low rates of penetration during the day shift may be associated with lower rates of attack from hackers, who often work in the evenings or nights. Similarly, higher rates of security breaches may, after detailed study, be found to have been caused by factors entirely outside the control of a particular employee.

In addition to monitoring performance, managers must ensure that all employees know that they are being monitored. Warning notices, pre-employment agreements, and yearly policy reviews can ensure that staff develop no unwarranted expectation of privacy about their work.[54]

One of the most effective supervisory practices an INFOSEC team leader, or any other manager, can use is *managing by walking around*.[55] Managers should set aside time every week to observe the conditions and to absorb the atmosphere of the working areas. Visiting team members and hearing about their specific job experiences, both

positive and negative, can only improve communications and motivation within the security team.

Not surprisingly, poor communication is a major enemy of effective team development. This poor communication can exist on multiple levels: among IT staff members, between staff members and IT management, and between IT and internal customers.[56]

Tools for enhancing communication include production control meetings, regular brief status review meetings, all-hands updates at which group members are required to give updates, and individual staff members to give technical updates, as well as establishing regular and thorough communication with internal customers and top management.

All members of the IT staff should be familiar with all lines of communication and with what behavior is expected of them in every scenario. Pushing communication awareness down through the organization appears to be an effective strategy.

63.2.4 Judgment and Adaptation. Management must not permit policies and procedures to keep the work from getting done. The comic strip *Dilbert*[57] has become popular largely because it caricatures managers who apply policy unintelligently; for example, in one real company known to the authors, managers decided to give the marketing department new laptops because of all the traveling they had to do. The managers then decided that to prevent theft, the laptops should be so securely fastened to the employees' desks that they could not be moved at all.

When security policies interfere with productivity, the correct solution is rarely black and white. Usually, neither dropping the policy nor enforcing it without change is appropriate. A hospital security administrator, for example, might note that a workstation in the emergency room is always logged on for the entire day, using the ID and password of the first person who logged on. Clearly, this violates the principle that everyone using the system must be positively identified and authenticated. Piggybacking on the first user damages the credibility of log files and makes it impossible to ascertain exactly which person is retrieving and modifying data at any time during the day. However, cracking down insensitively on the emergency room staff is a bad idea; chances are that the harried medical and support personnel are simply racing to get their work done saving lives. Logging off and on repeatedly is not a good method of identification and authentication in that environment. A reasonable security manager would listen to the employees, understand their point of view and their functional needs, and then explore technical alternatives to the usual ID-password technique. For example, the security manager might find that proximity cards or smart cards could meet the requirements at reasonable cost.[58]

Management must concern itself with safeguarding the resources under its jurisdiction. Just as investors seek both a high rate of return and reasonable safety for their investments, so must managers seek a high rate of return through the effective use of resources under their command and take adequate steps to protect the value of the resources.

The manager's function is essentially the management of resources: human resources and capital resources. In a computer operations environment, capital resources are represented by the investment in equipment and operating programs. Human resources are represented by the skills needed to operate and control both hardware and software facilities. Human capital resources are also represented by complete operating programs. Information is another form of resource, one that is often created as a product of data processing or concentrated at the computer facility, in order to utilize equipment resources better.

63.2.5 Management Failures. The most pervasive vulnerabilities in computer security are due to poor management. Efforts to contain or mitigate computer security exposures often fall short, or fail, because of management inadequacy, either of senior management or of operating management, or both. Sometimes the problem is simply inertia. Equally damaging is management lip service. In far too many instances, computer security is simply not taken seriously by senior management. Some of the most common management errors are listed next[59]:

- The belief that "if it hasn't been needed before, it probably will not be needed now." This is the default frame of mind, which is really self-insurance.
- Competition with other goals. Security measures use resources that are often needed for other activities. If security is regarded as an add-on burden rather than an integral part of the business process, it may be neglected or postponed.
- Lack of contribution to the bottom line. If security needs to be cost justified as an independent activity, it may be regarded as a target for elimination or reduction.
- Unwillingness properly to fund security activities. Far too often management is reluctant to commit all the resources needed for complete protection. A prime example is the lip-service contingency plan, published and display to the auditors, but never fully implemented.
- The explosive growth of the Internet and of the Web has increased the number of novice computer users into the hundreds of millions worldwide, cowering in fear at the latest hoax but cheerfully sending each other joke programs with embedded viruses. Some of these novices have taken on responsibilities for computer system management and have spread havoc within their organizations.
- The common occurrence of laptop theft has shown how vulnerable *any* kind of unencrypted removable media can be. Several companies have come under fire for not having a laptop encryption policy. This policy should be extended to any type of removable media (such as USB "thumb" drives) for certain types of information.
- Cell phones and PDAs are other technologies that need to be controlled. They are frequently connected to company computers, on the company network, to sync up with Microsoft Outlook email and calendars. These provide another potential point of entry for malicious software.
- Adults such as teachers and parents (many of them violating software copyrights without realizing they are breaking the law) have too often failed to teach children in their care how to resist the wiles of criminal hackers, virus writers, pornographers, and pedophiles. Some of these children have now grown up and become young managers who tolerate, encourage, or demand illegal acts by their employees.
- Some Websites are being managed by undertrained staff who know nothing about the years-old vulnerabilities they have left invitingly on their systems. These unfortunate people are stuck with inadequate resources and dismissive managers, who nevertheless blame them when the site is plastered with obscenities by teenagers with little or no conscience. Even worse, the Website may be used in a for-profit scheme by the ever-increasing legions of highly skilled professional hackers.
- In addition to the old vulnerabilities, these undertrained (or simply overworked) staff may also be introducing new vulnerabilities to their Websites by using new

tools or methodologies with no consideration of security. These new methodologies, such as the current AJAX (Asynchronous JavaScript and XML), may make use of old tools, but may do so in a new and insecure way.

- Some Web designers assume that users should trust mobile application programs (Java applets, ActiveX controls) whose origins are uncertain, whose documentation is unavailable, and whose actions may be pathological. Technically incompetent managers who take credit for their corporate Website may have no idea of what their Webmasters are doing—and failing to do—to protect their site.[60]
- Software makers, who ought to have known better, have blurred the distinction between document and program by adding automatic execution of macros to their word processors. Emailed Trojan horses are activated automatically when the message is opened. Bloated programs are routinely so full of bugs that consumers now think it is normal to pay money for a service release that fixes what never ought to have been released. Some managers with inadequate training take this situation as a given.
- Not necessarily a failure on just the part of management, software is often written without thought to security at all. Buffer overflows continue to be the "low-hanging fruit" for software hackers. Unfortunately, secure coding has been slow to catch on, but is showing signs of gaining traction today.[61] SANS has announced it will offer courses in secure coding, and there is an expanding list of books and papers on the subject. Perhaps as it becomes mainstream, universities will start requiring full courses in secure coding practices for computer programming degrees.
- We collectively continue to use Internet protocols devised decades ago and which have no provision for packet authentication. Criminals forge mail headers and packet headers with impunity, and use them for denial-of-service attacks and email spam. Network managers continue to avoid output filtering, which could reduce such attacks, because it is too low on their list of priorities, or because they have never considered their responsibility to other Internet users.
- Access control still relies largely on the outdated and ineffective use of passwords chosen by untrained users. Conveniently for criminal hackers, many users pick names of family members, people with whom they are having illicit romances, movie stars, pets, favorite sports teams, and the names of objects on their desks or visible from their windows. Some managers contribute to this situation by refusing to follow recommendations for better passwords and by refusing to consider better, or additional, authentication methods, such as tokens and biometrics.
- Some managers are loyal neither to colleagues nor to employers. The boom in firings and job-hopping has led to a shortsighted emphasis on the quarterly bottom line that makes investments in corporate security seem pointless.
- Security specialists still lack reliable data on network and computer intrusions. Although there are now some resources for information sharing, such as CERT/CC, many managers resist contributing to the knowledge base for fear that they and their enterprise will be embarrassed by their victimization.
- Social engineering is still one of the largest weaknesses of the security realm. Hackers really do not need to spend much time trying to break through defenses when so many users are perfectly willing to give them free access to go right through. While it is debatable whether training programs will ever be sufficient to overcome this, managers need to pay close attention to the "phishing" attacks that are now becoming more popular and more sophisticated.

Some organizations are taking steps to implement creative and memorable security awareness programs such as *Jeopardy*-style games, where workers compete to supply the right answers to security-related topics. Scavenger hunts on a company's Website are used to find the answers to 10 security-related questions.[62]

63.2.6 Risk Management. Part of the data processing management task is to protect information resources and to safeguard the human capital resources that are essential to the services provided. Top management must concern itself with adequate recognition of the risks and must be assured that protective measures relative to these risks are in effect.[63]

In addition to the problems caused by management failures discussed in Section 63.2.5, managers should be aware of other risks in their work such as:

- **Physical hazards.** The likelihood of threats, whether accidental or intentional, that can result in physical damage. Fire, water, power loss, explosions, vandalism, terrorism, and civil disorder are all within this category.
- **Equipment malfunction.** The possibility of failures in computers and supporting equipment, such as printers, disk drives, and air conditioners.
- **Software malfunction.** The likelihood of loss and failures caused by computer programs, including operating system software and application programs.
- **Human error.** The threat of disruption or loss due to accidental or intentional action or inaction by employees. Computer operators, programmers, maintenance engineers, and service personnel can all precipitate loss.
- **Misuse of data.** The capacity for intentional misuse of information or facilities by perpetrators of crime, such as fraud, espionage, misrepresentation, forgery, or theft of data or other assets controlled by the data.
- **Loss of data.** The intentional or unintentional loss of information through disruption of the physical media on which the data resides, or the corruption or erasure of the data.

In addition to classifying risks by category of threats, it is useful to analyze risks by the magnitude of potential loss, the probability of loss, and the frequency and permanence of occurrence. Although magnitude of loss can be expressed in terms of time or dollars, it is more practical to use dollar cost as a common basis for measurement. Quantifications should be based on reasonable and supportable estimates of the costs associated with the actual occurrences of adverse events.

Threats must be evaluated in terms of probability or occurrence. True risk is difficult or impossible to measure, but a reasonable priority of risks may be established by evaluating the likelihood of occurrence in conjunction with the magnitude of potential loss for each threat. An aggregation of consequential costs for each threat, over a common time period, and based on the likelihood of occurrence for each threat, can serve to prioritize seemingly diverse risks.

Three convenient groupings for permanency of damage are disasters, solid failures, and transient failures. Disasters are serious and lengthy disruptions, usually resulting in costly reconstruction of data, alternative off-premises processing, loss of business, and high cost. Solid failures are those that require shutting down part or all of the system in order to take corrective action. Costs of solid failures may range from simple inconvenience to substantial loss of business. Transient failures are defined as

temporary disruptions that do not recur regularly and therefore may be difficult to correct. Like solid failures, their costs can vary widely.

Quantification of risk is imprecise at best and varies greatly from one organization to another, and even among computer installations in the same organization. Nevertheless, quantification affords the means of ordering the relative importance of various threats and of substantiating the need for expenditure to counteract threats. In summary, even though probability estimates for specific events may represent intuitive feelings about likelihood, rather than actuarial knowledge, they can nonetheless serve as a basis for focusing discussion and planning.

An effective computer security program requires a balance of rationality and prudence. It also requires a continuing management commitment. Absolute security is an impossible dream unless one has unlimited resources—and even then, many responses would still have to be post hoc reactions to new threats, vulnerabilities, and attacks. Surprisingly, however, even on a modest budget, it is usually possible to achieve reasonable security. Many basic safeguards can be implemented for modest expenditures of time and effort. Many chapters in this *Handbook* discuss such basic safeguards.

63.3 LIABILITIES. As discussed, security managers focus on minimizing liability by a practice generally known as *risk management*. Risk management is the traditional model for INFOSEC; ideally, one determines risk by identifying the threats and vulnerabilities and then evaluating the associated costs and probabilities of each type of incident. The probabilities are difficult to define, and as a result, much of risk management is, in practice, an intuitive, nonquantitative process.

Security managers face many liabilities. Some of the possible negative consequences of inadequate security include:

- Loss of revenue
- Loss of reputation
- Loss of business partner confidence
- Loss of consumer confidence
- Loss of enterprise valuation
- Failure of the entire enterprise

Each of these types of loss also involves a loss of trust. Trust is easy to lose, especially in uncertain economic climates, and it is, unfortunately, harder to regain trust than to establish and maintain it. It is necessary for managers to understand that security and privacy are integral to the services and products offered by the enterprise. Security and privacy must apply to data from customers, business partners, employees, and every other individual or entity with whom the enterprise comes into contact.

63.3.1 Stakeholders. It is easy to think of only stockholders and customers when evaluating the potential costs of security breaches. However, it is useful to enumerate all of the people, and other entities, that are potentially affected by INFOSEC breaches. Such *stakeholders* can include:

- **Stockholders.** People and organizations owning stock in a privately held or publicly traded enterprise
- **Employees.** Managers and workers depending on an enterprise for their livelihood

- **Customers.** People and organizations depending on fulfillment of contractual obligations by the enterprise
- **Potential customers.** Those who might want to do business with the enterprise
- **Suppliers.** Those depending on the enterprise for acceptance of materials and services followed by payment as per contract
- **Data subjects.** People or other entities about whom an enterprise stores, manipulates, and reports data
- **Regulatory agencies and law enforcement.** People and organizations devoted to enforcing statutory regulations and laws
- **Users of other systems victimized by means of a compromised system.** Innocent bystanders who may be harmed, through no fault of their own, when a compromised system is used as an intermediary to launch attacks on an ultimate target

Managers must explore the potential consequences of specific types of security incidents with respect to all of the stakeholders. This wider analysis contributes significantly to effective risk management.

63.3.2 Due Diligence of Care. Due diligence of care in INFOSEC refers to the research and analysis carried out in establishing that risks have been minimized to an extent consistent with industry standards. Due diligence investigations, typically, are crucial in mergers and acquisitions, where unanticipated liabilities resulting can result in financial disaster and legal culpability.

Unfortunately, just as there is no sound basis for assertions about computer crime rates, there is not even the most rudimentary basis for asserting that any given level of security represents adequate care for information. What, precisely, is the right length for the asymmetric encryption key used in protecting confidential email sent across the Internet? Does failing to update antivirus signatures daily constitute a violation of due diligence, or is once a week good enough? Does due diligence in securing a manufacturing system require installation of an intrusion detection system? What about lacking an IDS in a hospital or a bank? Is a computer emergency response team a requirement to demonstrate due care in protecting information assets?

Attempts at defining security standards have not yet convinced more than a few enterprises to conform to their recommendations; the INFOSEC field is not yet at the point where quality assurance was a decade ago, when the International Standards Organization promulgated ISO 9000 certification, and it became widely used in manufacturing plants around the world.

Sensitivity to due diligence of care should, at a minimum, begin with consideration of legal and regulatory requirements for the protection of information. Contractual obligations with any and all stakeholders will determine the required degree of responsiveness to intrusion, and may help determine whether to cooperate with law enforcement authorities to investigate any breach of security.[64]

63.3.3 Downstream Liability. In recent years, a growing number of security experts and attorneys have predicted that the doctrine of downstream liability would become a significant factor in pushing management toward better security. Downstream liability is an application of the legal theory of contributory negligence as it applies to INFOSEC. In turn, contributory negligence refers to reckless endangerment of

others, "reckless" meaning without consideration of the consequences and endangerment resulting from putting others at risk of harm. The term "downstream" refers to the conventional model in which the source of data is viewed as being upstream of the recipient of those data; thus if someone compromises a university computer system and uses those computers to launch an attack on a bank, the bank is viewed as downstream from the university. Conceivably, the bank's attorneys could accuse the university administrators of negligence for allowing their computers to be compromised and therefore claim damages.

Keeping in mind that the authors are not attorneys, and that the information being discussed here is in no sense legal advice (for legal advice, consult an attorney), examples of what might be construed in a court of law as downstream liability include:

- Distributing virus-infected documents through email because all antivirus mechanisms have deliberately been turned off
- Allowing a malefactor to install *zombie* software on a poorly configured system and involving that compromised system in a distributed denial-of-service attack[65]
- Failing to install patches for a well-known and years-old vulnerability, thus allowing a criminal hacker to attack a third party via a root compromise of the poorly secured system[66]
- Allowing private information such as credit card numbers belonging to thousands of people to be stolen and distributed on the Internet, for use in credit card fraud
- Providing an unsecured email server that provides an open spam-relay point for junk email, with a forged REPLY-TO address, to flood millions of mailboxes with unwanted email, thereby causing thousands of bounce messages and angry accusations to clog the mail system of the innocent owner of the forged REPLY-TO address
- Having an employee who sends out thousands of fraudulent notices using the employer's email system to libel a competitor, causing depression of that competitor's sales and stock price
- Configuring a *honey-pot* system on the enterprise network to attract the attention of criminal hackers—who then turn around and use the honey pot to attack a third party

It is important to note that *honey pots*, while used as security measures to attract a hacker to a specific server for the purpose of gathering data on an attack, do in fact introduce security holes. The honey pot is a live server, with intentional vulnerabilities built in, connected to a company's network. Once penetrated, an attacker may use it as a platform from which to launch future attacks. Special safeguards must be designed to prevent this.

One possible safeguard is to prevent any traffic that leaves the honey pot to actually go anywhere, or to trap it on what might be termed a "honey logger," a server designed to simply catch all the traffic for the purpose of analyzing it.

As of this writing, there appear to have been no cases in which a plaintiff has successfully sued an enterprise for damages on the basis of downstream liability linked to inadequate security. However, in December 2000, FirstNet Online (Management) Limited filed a lawsuit in the Court of Sessions in Edinburgh, Scotland, against Nike, Inc., apparently an innocent victim. The incident began in June 2000, when someone hacked the DNS (Domain Name System) by filing incorrect data for resolution of the

nike.com domain with Network Solutions, Inc., which is charged with the responsibility of managing acquisition and retention of certain classes of domain names. All subsequent attempts to reach nike.com were redirected to the s11.org domain, an activist site devoted to fighting globalization. The redirection allegedly caused serious harm to the Web-hosting service for s11.org; according to the plaintiffs, FirstNET Online "experienced an 1800% increase of traffic over the 46 hours it took to correct the problem completely."[67] Blaming Nike administrators for failing to protect the password required for updating the DNS records at Network Solutions, the plaintiffs demanded compensation for the expenses incurred. Nike officials rejected such accusations and blamed Network Solutions for allowing the redirection of Internet traffic to the s11.org site. Readers will want to search the Web for later developments in this interesting case.

63.3.4 Audits. As managers attempt to manage security according to nebulous principles of due care, they must remember that in many enterprises, IT and INFOSEC departments are seen as being equivalent to the police. Many people in these departments have privileged access to other people's secrets; for example, they can read all email, and can often tell exactly what Websites' employees have been visiting. In some environments, network monitors allow administrators and security personnel to activate keystroke monitoring and to view in real time the appearance of any given terminal or workstation on the network. Because of the enormous power of these people, managers must ensure that their levels of access are not abused.

Information systems auditors are responsible for keeping the technical and security staff honest. In some larger enterprises, there is a department of internal audit whose director reports to the same level as other officers, such as the CEO and CFO. In other cases, third-party auditors monitor adherence to policy, standards, and procedures.[68]

Security audits are often feared, and with good reason. In general, many audits are performed with the focus of finding fault. This approach is counterproductive and unnecessarily stressful. The point of an audit should not be to find the maximum number of offenses but rather to assess the level of compliance. The most important outcome of an audit is finding areas of potential improvement. A nonadversarial style facilitates positive results; thus, an auditing team should provide constructive suggestions for improvement wherever possible.

An auditing team also needs to listen to those being audited. Although a particular network may not be configured to full compliance with security policy, there may be reasons for this apparent failure. For example, suppose a specific patch was required by security policy to have been installed on a Solaris server, but a particular Solaris server did not comply. An auditor might report a failure without further investigation even though, in reality, the patch might have been installed but found to be incompatible with a critical application. Just like managers, audit teams need to see the big picture, not just the rules and procedures. The main point is that all audits must be done with the needs of the business in mind. It does no good to have a process that gets in the way of the business at hand. Instead, to enable the business, alternative measures must be found. That is the job of the auditor—to help ensure compliance and to provide another set of eyes that can see to it that all needs are taken care of.

To prove due diligence, auditing is a must. Generally, third-party audits create more trust, as a third party has no bias about the findings. However, it is important to ensure that a third-party audit firm maintains its independence. It can be risky to use a small firm that bases more than a modest fraction of its revenue on income from any one client. However, it may be undesirable to have the same large firm perform accounting or auditing functions while also serving as security consultants.

63.4 COMPUTER MANAGEMENT FUNCTIONS. The job of the manager is to provide leadership, and it must be provided in an organized and creative manner. Management is dynamic, not static, and the manager must deal with change—change in the organizational environment, change in people, and change in the methods of management. Unfortunately, the laws of inertia apply to management as to physics, and change will often be resisted. The challenge to the manager is to manage change.

Even a successful security program will usually result in higher costs, as well as in changes in the organizational structure and in the working environment. To compound the problems, the benefits of the security program are not highly visible. Security is essentially preventive and is often regarded as capital expenses and overhead of questionable value. Too often, upper managers dismiss proposals for better security by pointing out that there have been no breaches of security in the last year; thus the more successful the security program, the less evidence there is to support its usefulness. To counteract this circular and destructive reasoning, security managers should compile evidence of the numbers and types of attempted penetrations identified by intrusion detection systems[69] and attempted malfeasance by authorized users identified by system and application logging.[70]

The tools of the manager include planning, organizing, integrating, and controlling. These are not independent activities that can be completely separated; rather they represent a matrix. Taken as an interrelated process, they achieve balance and direction.

63.4.1 Planning for Computer Security. The planning function of management includes the determination of objectives, policies, priorities, schedules, standards, and strategy.

It is important to define the scope and purpose of a computer security program in terms of objectives. There should be a clear statement of results to be achieved within a given period of time. Security objectives must be balanced with other organizational objectives, because conflicts may arise. As an example, the need for controlled access will naturally conflict with desires for user flexibility and convenience. Objectives should be imaginative and responsive to change and conflict.

Planning for computer security requires the participation of top management, so that security objectives can be reconciled with general organizational objectives and with financial priorities. It is also necessary to coordinate security activities between data processing and all the other areas of the organization. Auditing, insurance, legal, financial, and other groups are affected by, and should contribute to, a computer security program.[71]

The objectives, policies, schedules, and standards that result from the planning process need to be communicated throughout the organization. While this is generally true for all planning, it is especially crucial to the success of a security program, which may conflict with existing corporate culture.[72] Finally, feedback is essential to permit recognition of failures and departures from plans. Only by monitoring results is it possible to take corrective action or to readjust the objectives, policies, plans, schedules, and standards to the practicalities of the real world.

63.4.2 Organizing. Organization is the process of marshaling resources, grouping activities and responsibilities, and establishing relationships that will enable people to work together most effectively in determining and accomplishing the objectives of an enterprise.

The important elements of organization for computer security include:

- Obtaining resources of personnel, money, and facilities adequate to accomplish the assigned mission
- Fitting responsibly into the organizational pattern
- Assigning responsibility and authority to individuals
- Formulating supporting methods and procedures
- Measuring organizational effectiveness

Traditionally, the organizing function is concerned with grouping activities into manageable components and grouping human resources into logical relationships to accomplish the desired results. It would be unusual to find an enterprise that was designed with an optimum structure to achieve security at the expense of other goals. Security is not an independent activity, nor does it represent the primary goal of a data processing organizational element. Data processing departments are usually organized into units that reflect the nature of the work performed, or user relationships, or some other structure designed to achieve an adequate service level for the end user. Security measures quite often conflict with the service objectives. That conflict is all the more reason why security must be managed in order to be effective. It is also motivation for creating a separate functional activity, particularly in large organizations, so that INFOSEC can be administered independently of competing activities.

63.4.3 Integrating. Computer security is frequently an afterthought to the organization of the data processing function. Security is also perceived as a passive activity. As a result, responsibility for security is often assigned haphazardly. Security should not merely be superimposed but, instead, should be carefully fitted into the organizational structure. Some important considerations are:

- Accountability for specific security tasks should be included in formal job definitions of every job level.
- Training programs should include a complete review of security objectives and policies, as well as details relating to assigned security tasks.
- Supervisory and management personnel should be assigned responsibility for both performance and attitude of staff with respect to security.
- Certain line or staff positions could include responsibility for overall security, or a grouping of security tasks, or measurement and monitoring of security. It may be convenient to combine security and control of data processing into a single function. Monitoring of security effectiveness is sometimes assigned to the data processing auditor.
- Primary responsibility for INFOSEC should not be assigned to those with inherently conflicting priorities; for example, it is a bad practice to assign major responsibility for security to the chief of operations, the director of software development, or the vice president of finance.
- Information systems security should be coordinated and reviewed with security specialists in other areas of the organization. In a manufacturing concern the plant security personnel, or in a bank the bank security officer, should be consulted. It

may be desirable to assign monitoring responsibility for security housekeeping tasks to these specialists.

- Whenever possible, policies should be translated into written guidelines and procedures to provide the detailed requirements for each task. Standards for performance of each activity must be formulated and applied.

Initiating and monitoring the accomplishment of objectives according to established plans requires skillful leadership. The integration function has a direct bearing on the success of a security program. Without effective leadership, security can become a farce. With leadership, a security program can overcome its basic passivity and truly enhance the utility of information systems. Management must embrace and inspire the concept of security as an everyday fact of life. But beyond transmitting a positive attitude of acceptance, management must also be concerned with testing and improving security. The security program is dynamic rather than static. Management's role must be one of continued concern in order to identify, adopt, and adapt better methods to accomplish corporate security objectives.

63.4.4 Controlling. The final process in the management cycle is measuring and controlling. It is the function that is concerned with achieving cost-effective results and includes establishing standards, improving methods, examining results, and adjusting the organizational mechanism for corrective action.

The basis for effective control is the use of consistent techniques for measurement and the application of standards for comparison and interpretation. Results should be analyzed promptly, as feedback to effect corrective action. Control systems need not be burdensome or elaborate, but they should be consistent, and they should allow for flexibility and adjustment. Management action is the end product of control. Mistakes in original objectives and plans are forgivable, but failure to recognize and react to mistakes can only compound the problem, add to costs, and undermine the effective use of valuable resources.

63.5 SECURITY ADMINISTRATION. In large organizations, an independent security administration function is often the most effective method for accomplishing the overall objective of improved INFOSEC.

63.5.1 Staffing the Security Function. The importance of INFOSEC has given rise to a new management specialty, consisting of professionals involved in the planning and administration of protection for the integrity and security of automated information assets. Titles such as INFOSEC administrator, computer security manager, information systems security officer, and chief INFOSEC officer have been used to describe these roles.

Coordination of the INFOSEC function requires a combination of managerial and technical talents. The successful administrator must be a superior communicator capable of selling the concept of security and maintaining security awareness at all levels of the organization. Sufficient technical knowledge is important, so that the INFOSEC administrator can evaluate and initiate appropriate technological solutions to meet corporate INFOSEC policies and to counteract threats. While technical skill, in the form of a data processing background, is important, a broader range of capabilities is needed for maximum effectiveness. The ideal security administrator should possess the

ability to communicate with all levels of management and should have good knowledge of related functions such as auditing, internal control, and general security. It is also important to have some knowledge of the industry within which the organization operates.

The administration of INFOSEC can be centralized or decentralized, depending on the needs of the organization. Where multiple data center locations are involved, the decentralized approach may be more appropriate to accomplish the details of administration. However, it is important to have one focal point for overall coordination of INFOSEC policy. It is also essential to understand that responsibility for INFOSEC rests with all members of the organization, and not just the security personnel. Security is a shared responsibility, and this concept must be widely promoted by the security administrator and strongly backed by senior management.

63.5.2 Authority and Responsibility. As in all management areas, INFOSEC administrators (ISAs) need not only specific responsibilities but also the authority to carry out their duties. Those duties are listed next:

- **Establish policy statements and guidelines for information protection.** Although policy is the primary responsibility of senior management, it is appropriate for the ISA to participate in the delineation of a formal policy statement covering this important organizational goal and to prepare appropriate guidelines.
- **Identify vulnerabilities and risks.** The ISA serves as a consultant and coordinator in the process of risk analysis. The sensitivity of data resources must be decided by senior management, but with full consensus and agreement by all affected sections of the organization. The ISA has a special responsibility to identify specific risks that affect the automated data resources. The ISA should then coordinate the process of quantifying or otherwise prioritizing the value of the vulnerable data, in order to establish a basis for selection and economic justification of protective measures.
- **Recommend protective measures.** Major responsibility for identification of economic solutions to INFOSEC vulnerabilities is usually assigned to the ISA. Requiring a combination of technical knowledge and management analysis, this process entails the evaluation of protective solutions for technological, operational, and economic effectiveness. Appropriate recommendations must be coordinated with other affected sections of the organization, including audit, data processing operations, software development, legal counsel, human resources, facilities security, public relations, and others. Implementation plans must also be developed, and there must be a management commitment to the implementation.
- **Control the implementation of protective measures.** Whether the final administration of the day-to-day security procedures is centralized or decentralized, the coordination and control of implementation for major protective measures should be centralized. A prerequisite for implementation is the development of standards for INFOSEC to ensure consistency in the application of protection. Important areas for standardization are security design for application systems, programming development, data sensitivity criteria, database access, and program maintenance. In general, security standards cover the entire systems life cycle.

- **Measure effectiveness of security precautions.** Feedback is essential to assess how effectively policies are being followed. Since the nature of INFOSEC is defensive, the measures adopted can easily fall into disuse unless there is ongoing confirmation of effectiveness. The ISA should have primary responsibility to conduct security audits for operational systems as well as for systems under development. Backup protection and disaster-recovery procedures are especially sensitive plans that must be tested periodically. Monitoring of variances in security procedures is also important, and is best controlled through the ISA function. In many organizations these activities are coordinated with the auditors. Finally, the ISA should provide senior management with reports on the effectiveness of security policy, with identification of weaknesses, and with recommendations for improvement.
- **Promote security awareness and security education.** Another important area of responsibility for the ISA is security education and awareness. The concept of security must be actively communicated to all members of the staff to maintain awareness of its importance. An effective program should achieve a workable balance between security and the utility of computer resources.
- **Ensure security awareness across cultures.** In today's global economy, many companies are located in various geographic regions. As a result, security professionals must understand the role that culture plays in the global enterprise. Different cultural perspectives exist regarding security, and these must be taken into consideration when developing a security-awareness program. Understanding the impact of culture on business operations will be very beneficial for the organization. For example, how should global enterprises create cross-border awareness strategies? How might the laws and regulations in various countries help or inhibit the creation of an awareness program?

  Although English is still a widely used language of business, U.S. security professionals are well advised to understand the cultural nuances of their foreign counterparts in order to protect their organization against increasing threats. The growing economic power of the European Union, China, and India, should prompt U.S. professionals to begin expanding their language competencies to match those of their global competitors.

63.5.3 Professional Accreditation and Education. Professional accreditation can provide managers with a basis of assurance that a common body of knowledge is applied to the requirements of INFOSEC in their enterprise. There are several professional societies and organizations that offer specialized certification to INFOSEC personnel. Such certification is voluntary, since there are no specific licensing requirements for INFOSEC practitioners. However, certification does signify serious professional intent on the part of the individual to acquire and maintain the needed knowledge and skills. In most cases, the certification process requires a combination of education, experience, and knowledge; knowledge is evaluated through written examinations. Once achieved, a professional certificate requires continuing education in the field to maintain its validity.

Unfortunately, some managers fear that after they provide training or support advanced education, their employees will leave and use their new skills to benefit some other organization—perhaps even a competitor. In reality, some employees are more likely to walk out if the investment is *not* made. Numerous studies have shown that if

employees feel valued, they will be more willing to stay. For example, a summary of changing contributions to employee satisfaction specifically states that:

- Although money is clearly an important consideration, cash is not the primary factor that keeps people in their current job or attracts them away to a new job.
- Opportunity to grow and learn at work is emerging as a primary determinant of attracting and retaining employees. …
- Although employees consistently indicate that education and training are key aspects of their willingness to stay at or leave their current job, employees also typically rank the quality of their employer's education and training function as low.[73]

Training need not be expensive, excessively time consuming, or difficult.[74] Online security training can be taken anytime and is self-paced. Training videos are also an excellent tool for stimulating employee knowledge and interest. The simplest way to locate such resources is to search on the Web with a good search engine using keywords such as "security training." A particularly helpful site is CCCure.org, which includes a wealth of self-study materials as well as links to many vendors.[75]

Another alternative is to have the employer use obsolete equipment for training, equipment that would otherwise lie idle in a storage area. Even small departments may have old PCs lying around dormant, as well as hubs and possibly router/switches. These can be used to create a network for testing and learning at little cost. Companies that do not have spare equipment can buy used equipment for less than the cost of a one-week intensive course. For an employee who wants to learn the basics of firewalls, Linux is a cheap (and even free) operating system that has just recently started supporting firewalls. Setting up a Linux firewall requires no expensive software or appliances. It can provide an effective way of encouraging employees to learn the practical details of configuring firewalls—a valuable skill set in any security department—while rewarding loyal employees by showing confidence in their commitment to learning, and to their continued employment in the enterprise.

Once the firewall is set up, other employees can use the testing and training network to learn about penetration testing, intrusion detection, and other security elements. Different operating systems can be installed, with various applications, all for a minimal investment. As technology, techniques, and tools change, this training network will be valuable in keeping skills up to date.

External, college-level undergraduate and graduate education also provides resources for employees and managers. Universities and colleges are increasingly recognizing the value and validity of educational programs centering on information assurance, both at the technical and managerial levels. Some offer noncredit courses or certificate courses in addition to degree programs. Some offer online programs that support the needs of working adults. Managers can take advantage of these opportunities to improve the knowledge base within their organizations and to increase employee loyalty and retention by supporting employees who want to further their careers in information assurance. Some programs include case studies that may, with permission, center on fieldwork within the students' own work environments, resulting in extensive analysis and recommendations of immense value to their employers.[76]

63.6 CONCLUDING REMARKS. Society needs the insights of information assurance specialists. Managers with daily opportunities to think about the strategic

implications of vulnerabilities, exploits, and, above all, about the business processes for which they are responsible are in an excellent position to contribute not only to the success of their own enterprises but also to the future of their societies. As Adam Shostack and Andrew Stewart have written in their future-pointing treatise, it is time for a new school of INFOSEC—one that involves:

> Learning from other professions, such as economics and psychology, to unlock the problems that stymie the security field. …
>
> Sharing objective data and analysis widely. …
>
> The embrace of the scientific method for solving important security problems. Analyzing real-world outcomes is the best way for information security to become a mature discipline.[77]

Readers should embrace every opportunity to bring their experience and wisdom to a wider audience, and by speaking and writing they should share their insights with each other and with the general public.

63.7 FURTHER READING

Calder, A., and S. Watkins. *IT Governance: An International Guide to Data Security and ISO27001/ISO27002*, 5th ed. Kogan Page, 2012.

Goddard, J., and T. Eccles. *Uncommon Sense, Common Nonsense: Why Some Organisations Consistently Outperform Others*. Profile Books, 2013.

Kouns, J., and B. L. Kouns. *The Chief Information Security Officer*. IT Governance Publishing, 2011.

Laplante, P., and T. Costello. *CIO Wisdom II: More Best Practices*. Prentice-Hall, 2005.

Mellado, D., L. E. Sanchez, E. Fernandez-Medina, and M. Piattini, eds. *IT Security Governance Innovations: Theory and Research*. IGI Global, 2012.

Moeller, R. R. *Executive's Guide to IT Governance: Improving Systems Processes with Service Management, COBIT, and ITIL*. Wiley, 2013.

Patterson, T. *Mapping Security*. Prentice-Hall, 2005.

Rees, M., and M. Kennedy. *Fraud and the Human Factor: Fighting the Greatest Threat to Today's Corporation*. Bloomsbury Information Ltd, 2013.

Solms, S. H., and R. Solms. *Information Security Governance.* Springer, 2008.

Swiderski, F., and W. Synder. *Threat Modeling*. Microsoft, 2004.

Talbot, J., and M. Jakeman. *Security Risk Management Body of Knowledge*, 2nd ed. Wiley, 2009.

Wheeler, E. *Security Risk Management: Building an Information Security Risk Management Program from the Ground Up*. Syngress, 2011.

Wibbeke, E. S., and S. McArthur. *Global Business Leadership*, 2nd ed. Routledge, 2013.

Wood, C. C. *Information Security Roles & Responsibilities Made Easy*, 2nd ed. Information Shield, 2005.

63.8 NOTES

1. Because this entire *Handbook* in a sense provides the underlying details for the management of information assurance, most cross-references to other chapters have been relegated to endnotes to avoid cluttering the text.
2. See Chapters 60 and 62 in this *Handbook.*

3. For an extensive discussion of the role of the CISO, see Chapter 65 in this *Handbook.*
4. California State Legislature (2002) Senate Bill SB 1386, http://info.sen.ca.gov/pub/01-02/bill/sen/sb_1351-1400/sb_1386_bill_20020926_chaptered.html
5. G. Lin, "The Tao Perspective," in D. Lane, ed., *CIO Wisdom* (Upper Saddle River, NJ: Prentice-Hall, 2003), Chapter 4.
6. NIST, "Information Security System in the Systems Development Lifecycle," 2004 [Brochure], http://csrc.nist.gov/groups/SMA/sdlc/index.html
7. M. E. Kabay, "The Net Present Value of Information Security: A Paradigm Shift for INFOSEC and E-commerce," vol. 17 (2006), www.mekabay.com/infosecmgmt/npvsec.pdf
8. J. Bassett and D. Rothman, *A Seat at the Table for CEOs and CSOs: Driving Profits, Corporate Performance & Business Agility* (Bloomington, IN: AuthorHouse, 2007), p. 111.
9. This section is based on the paper by M. E. Kabay, "The VA Data Insecurity Saga," 2008, www.mekabay.com/infosecmgmt/vasaga.pdf
10. J. Brodkin, "Rating Apologies: Deep Regrets, from TJX to ChoicePoint, about Data Leaks," *Network World,* March 14, 2007, www.networkworld.com/news/2007/031407-wider-net-apologies-letters.html
11. G. Gross, "U.S. Agency Loses Data Containing 26 Million IDs," *Network World,* May 22, 2006, www.networkworld.com/news/2006/052206-us-agency-loses-veterans-data.html
12. G. Gross, "Lawmaker Calls on VA Head to Resign after Data Theft," *Network World,* May 25, 2006, www.networkworld.com/news/2006/052506-lawmaker-calls-on-va-head.html
13. "FirstGov.gov," "Latest Information on Veterans Affairs Data Security," *Military.com | Money* (n.d.) www.military.com/money/personal-finance/credit-debt-management/latest-information-on-veterans-affairs-data-security.html
14. Using keyword "data" in the search field at http://www1.va.gov/opa/pressrel/index.cfm provides a reasonable chronology
15. U.S. Department of Veterans Affairs, "Secretary Nicholson Provides Update on Stolen Data Incident: Data Matching with Department of Defense Providing New Details," June 6, 2006, www.va.gov/opa/pressrel/pressrelease.cfm?id=1134
16. Associated Press, "Data on 2.2M Active Troops Stolen from VA," *USA TODAY,* June 7, 2006, www.usatoday.com/news/washington/2006-06-06-veterans-data_x.htm
17. U.S. Department of Veterans Affairs, "Directive by the Secretary of Veterans Affairs R. James Nicholson to All VA Supervisors on Information Security," May 26, 2006, www.va.gov/opa/pressrel/pressrelease.cfm?id=1128
18. U.S. Department of Veterans Affairs, "VA Secretary Inserts New Leadership in Policy & Planning Office," May 30, 2006, www.va.gov/opa/pressrel/pressrelease.cfm?id=1129
19. U.S. Government Accountability Office, "Homeland Security: Guidance and Standards are Needed for Measuring the Effectiveness of Agencies' Facility Protection Efforts," Report GAO-06-612, May 2006, www.gao.gov/cgi-bin/getrpt?GAO-06-612

20. U.S. Department of Veterans Affairs, "VA Secretary Inserts New Leadership in Policy & Planning Office."
21. U.S. Government Accountability Office, "Veterans Affairs: Leadership Needed to Address Information Security Weaknesses and Privacy Issues," Report GAO-06-866T, June 14, 2006, www.gao.gov/cgi-bin/getrpt?GAO-06-866T
22. U.S. Government Accountability Office, "Information Security: Leadership Needed to Address Weaknesses and Privacy Issues at Veterans Affairs," Highlights of GAO-06-897T, 2006, www.gao.gov/highlights/d06897thigh.pdf
23. U.S. Government Accountability Office, "Information Technology: VA and DOD Face Challenges in Completing Key Efforts," Highlights of GAO-06-905T, 2006, www.gao.gov/highlights/d06905thigh.pdf
24. Department of Veterans Affairs Office of Inspector General, "Review of Issues Related to the Loss of VA Information Involving the Identity of Millions of Veterans," Report No. 06-02238-163, July 11, 2006, www.va.gov/oig/pubs/VAOIG-06-02238-163.pdf
25. D. Pulliam, "Report: VA Treated Data Breach with 'Indifference'," Government Executive Website, July 11, 2006, www.govexec.com/technology/2006/07/report-va-treated-data-breach-with-indifference/22234
26. "Personal Data for 38,000 Veterans Missing, VA Says," *USA TODAY,* August 7, 2006, www.usatoday.com/tech/news/computersecurity/infotheft/2006-08-07-veterans-data_x.htm
27. G. Gross, "VA to Spend $3.7M on Encryption Tools: The Move Follows the Theft of a VA Laptop in May," *Computerworld,* August 14, 2006, www.computerworld.com/s/article/9002447/VA_to_spend_3.7M_on_encryption_tools
28. R. McMillan, "Update: Unisys Contractor Arrested in VA Theft: Investigators Do Not Believe 21-Year-Old Suspect Sought Agency's Data," *InfoWorld,* September 18, 2006, www.infoworld.com/d/security-central/update-unisys-contractor-arrested-in-va-theft-825
29. H. A. Waxman, "Committee Report Finds Data Breaches Throughout Federal Government," Committee on Oversight and Government Reform, 110th Congress, October 13, 2006, http://oversight-archive.waxman.house.gov/story.asp?ID=1127
30. R. Maze, "VA Reports Two More Data Security Lapses," *Federal Times,* November 3, 2006 (no longer available online).
31. A. Broache, "Hard Drive Vanishes from VA Facility," *C|Net News,* February 5, 2007, http://cnet.news.com/2100-1029_3-6156386.html
32. G. Keizer, "Lost VA Hard Drive May Have Held 1.8M IDs: Initially, the Agency Said Just 50,000 Were Potentially Affected," *Computerworld,* February 13, 2007, www.computerworld.com/s/article/9011218/Lost_VA_hard_drive_may_have_held_1.8M_IDs
33. M. E. Kabay, "PIIssed Off Yet?" *Network World | Security Strategies,* June 12, 2007, www.networkworld.com/newsletters/sec/2007/0611sec1.html
34. U.S. Government Accountability Office, "Information Security: Veterans Affairs Needs to Address Long-Standing Weaknesses," Report GAO-07-532T, February 28, 2007, www.gao.gov/new.items/d07532t.pdf
35. G. Gross, "VA to Spend $3.7M on Encryption Tools."
36. J. Vijayan, "OMB Sets 120-Day Deadline for Fed Breach-Notification Plans: Agencies Have the Summer To Develop and Implement First Phases of Policies,"

Computerworld, May 29, 2007, www.computerworld.com/s/article/9021544/OMB_sets_120_day_deadline_for_Fed_breach_notification_plans

37. Chronology of Data Breaches: Security Breaches 2005–Present," Privacy Rights Clearinghouse (2013-07-14), www.privacyrights.org/data-breach
38. See Chapters 1 and 3.
39. See Chapter 57.
40. See Chapters 56, 58, and 59.
41. See Chapter 60.
42. See Chapters 22 and 23.
43. See Chapters 24 through 37.
44. See Chapters 43 through 50.
45. See Chapters 51 through 55.
46. See Chapter 45 for more on employment practices and policies.
47. Policy issues are discussed in detail in Chapters 44, 45, 47, 48, 49, and 50 of this *Handbook*.
48. M. K. Pratt, "Employee Security Training: Beyond Posters," *Computerworld*, April 17, 2006, www.computerworld.com/s/article/110494/Employee_Security_Training_Beyond_Posters
49. G. G. Gross, "Study: Human Error Causes Most Security Breaches," *Computerworld*, March 18, 2003, www.computerworld.com/s/article/79485/Study_Human_error_causes_most_security_breaches
50. E. Dickson, "Could the Arrests in the Stop and Shop Data Breach Indicate a Tie to Armenian Mobsters?" Fraud, Phishing and Financial Misdeeds Blog (Feb 28, 2007), http://fraudwar.blogspot.com/2007/02/could-arrest-in-stop-and-shop-data.html
51. See Chapter 50 for more information about using social psychology to motivate compliance with security policies.
52. The Honeynet Project: Challenges, www.honeynet.org/challenges
53. For more about log files, see Chapters 52 and 53; on postmortem analysis, see Chapter 56.
54. See Chapters 49 and 50.
55. T. J. Peters and R. H. Waterman, *In Search of Excellence: Lessons from America's Best-Run Companies* (New York: HarperCollins, 1983). See also H. K. Jones, "Does MBWA (Management By Wandering Around) Still Work?" Self-Growth.com, 2007, www.selfgrowth.com/articles/Jones6.html
56. B. J. Fox, "Communications: Communication Excellence in IT Management," in D. Lane, ed., *CIO Wisdom,* (Upper Saddle River, NJ: Prentice-Hall, 2003), Chapter 5.
57. S. Adams, "Official Dilbert Website" (2008), www.dilbert.com
58. See Chapter 28.
59. M. E. Kabay, "A Rant about InfoSec: A Security Veteran in a Bad Mood Dumps on Everyone" (2004,; www.mekabay.com/opinion/rant.pdf
60. See Chapter 30.
61. See Chapter 38.
62. For more security-awareness ideas, see Chapter 49.

63. Chapter 62 of this *Handbook* discusses risk management in more detail.
64. See Chapter 61.
65. See Chapter 18.
66. See Chapters 15 and 40.
67. R. J. McGillivray and S. C. Lieske, "Webjacking," *The Computer & Internet Lawyer* 18, no. 7 (July 2001):1, www.hackerzvoice.net/ouah/webhijack.html
68. See Chapter 54.
69. See Chapter 27.
70. See Chapters 53 and 54.
71. See Chapter 56 for a more extensive discussion of these principles in connection with the computer emergency response team.
72. See Chapter 50 for a discussion of the psychology of implementing security policies.
73. L. J. Basai, "Employee Retention." Learning in the New Economy," *LineZine* (Summer 2000), www.linezine.com/3/themes/hardtalk.htm
74. M. R. Farnum, "Security Awareness Training Does Not Have to Be Hard," Hitting the Security Nerve—*Computerworld* Blogs (2006), http://blogs.computerworld.com/node/4175
75. www.cccure.org
76. For more details of United States–based certification and education in the security field, see Chapters 74 and 75.
77. A. Shostack and A. Stewart, *The New School of Information Security* (Upper Saddle River, NJ: Addison-Wesley, 2008), p. xiv.

CHAPTER 64

U.S. LEGAL AND REGULATORY SECURITY ISSUES

Timothy Virtue

64.1 INTRODUCTION 64·1

64.2 SARBANES-OXLEY ACT OF 2002 64·2

64.2.1 Section 404 of SOX 64·4
64.2.2 Management Perspectives on SOX 64·5

64.3 GRAMM-LEACH-BLILEY ACT 64·6

64.3.1 Applicability 64·6
64.3.2 Enforcement 64·7
64.3.3 Consumers and Customers 64·8
64.3.4 Compliance 64·8
64.3.5 Privacy Notices 64·9
64.3.6 GLBA Safeguards Rule 64·9
64.3.7 Flexibility 64·10

64.4 EXAMINATION PROCEDURES TO EVALUATE COMPLIANCE WITH GUIDELINES FOR SAFEGUARDING CUSTOMER INFORMATION 64·14

64.5 CONCLUDING REMARKS 64·14

64.6 FURTHER READING 64·14

64.7 NOTES 64·15

64.1 INTRODUCTION. Since the previous edition of this *Handbook* was published in 2009 there have been significant technical, political, economic, legal, and regulatory changes that have complicated efforts to achieve compliance. In addition, the overall global economic downturn has made compliance more difficult.

With the global and U.S. economic downturn, many organizations were challenged to maintain or increase their compliance programs with fewer resources. In periods of economic decline, fraud and other criminal activities have a tendency to increase, requiring security leaders to focus on both security and regulatory issues simultaneously.

The rising use of social media, mobile devices, and cloud computing has complicated the application of security controls and strategies and increased regulatory and legal challenges. Security teams must apply flexible, innovative, and adaptive security solutions.

Many organizations are restructuring independent and isolated operational units (sometimes described as *silos*) and focusing on coordinated strategic risk management for the enterprise. Such changes support business values, achieve compliance, and reduce risk.

Another recurring theme of today's regulatory environment is accountability for, and protection of, sensitive or private information. Specific regulatory requirements and how an organization addresses the requirements will vary, based on industry, organizational structure, and internal processes. However, an effective enterprise-wide regulatory-governance model should incorporate two common elements:

1. The senior management team must create a compliance-driven culture by actively exercising leadership by example.
2. The team must actively support this culture of compliance by enabling computer security strategies, tactics, and technologies as discussed throughout this *Handbook*.

Compliance with regulatory requirements must include, but is not limited to, these computer security fundamentals:

- Program management
  - Policy and procedure design and implementation
- Risk assessment and management
  - Prevention
  - Detection
  - Response
- Implementation of safeguards
  - Administrative controls
  - Physical controls
  - Technical controls
- Awareness and training

Organizations must also keep in mind that vendor due diligence and management may play a more critical role as they adopt social media, mobile, and cloud-computing technologies within the organization. The same levels of management and oversight still need to exist; however, now they need to be applied outward.

Although there are a number of regulations that today's business leaders may be accountable for, these requirements vary from industry to industry. This chapter focuses on two important U.S. regulations: the Sarbanes-Oxley and the Gramm-Leach-Bliley Acts. These particular acts require a significant level of management involvement in order to validate that the organization adequately addresses the regulatory requirements.

64.2 SARBANES-OXLEY ACT OF 2002. The *Sarbanes-Oxley Act of 2002* (SOX), also sometimes referred to as the *Public Company Accounting Reform and Investor Protection Act of 2002*, comprises 11 titles and is administered by the Securities and Exchange Commission (SEC). The main goal of the legislation is to protect the public and shareholders from accounting errors and fraudulent practices by requiring that companies monitor and record any access to sensitive financial data. They must provide reports on their controls, and must respond in a timely and appropriate manner should an incident occur.[1]

EXHIBIT 64.1 Eleven Titles of SOX with Descriptions

Title #	Name	Description
I	Public Company Accounting Oversight Board	Establishment of a private nonprofit board. Duties of the board include establishing auditing, quality control, and independence standards for public accounting firms.
II	Auditor Independence	Prohibits certified public accountants from simultaneously performing audit and nonaudit services.
III	Corporate Responsibility	Requires that the Securities and Exchange Commission (SEC) order the national security exchanges and associations to bar the listing of an issuer that does not comply with the specific requirements set forth in Title III.
IV	Enhanced Financial Disclosures	Sets forth various requirements to validate the accuracy of financial statements and supporting disclosures.
V	Analyst Conflicts of Interest	Defines a code of conduct for security analysts and includes a measure to help restore confidence in their reporting.
VI	Commission Resources and Authority	Authorizes the SEC to hire staff to provide additional oversight of auditors and audit services. Also outlines the conditions a person can be prohibited from practicing as a broker, advisor, or dealer.
VII	Studies and Reports	Provides for research to be conducted for the enforcement of actions against violations by SEC registrants and auditors.
VIII	Corporate and Criminal Fraud Accountability	Defines penalties for the destruction of audit reports and the willful destruction, manipulation, or falsification of documents in federal investigations and bankruptcy proceedings.
IX	White Collar Crime Penalty Enhancements	Increases the penalties for white-collar crimes and requires that both the CEO and CFO of a corporation certify periodic financial statements.
X	Corporate Tax Returns	Requires that the CEO sign the corporation's federal income tax return.
XI	Corporate Fraud and Accountability	Establishes penalties for persons who manipulate or destroy documents or impede an official proceeding.

Exhibit 64.1 offers each title within SOX with a brief description added by the author.

All of the individual components from Exhibit 64.1 that specify the requirements for SOX must have the support and involvement of an organization's senior leadership team.

No single component of SOX is more important than any other component. However, the remaining discussion focuses on the part of the Act that affects computer security the most: Section 404.

The passage of SOX in 2002 led the way for an important change in how business leaders in today's public companies look at information security. This change has caused significant impact on today's information systems, the security that supports them, and management's perception of how the systems affect financial reporting.

Although SOX was primarily intended to enhance the integrity of financial reporting, it has broader implications. Specifically, the integrity of financial reporting can be accomplished only by ensuring the integrity of the source information used to complete financial reporting requirements. Management must therefore support the design, implementation, and continuous management of strong controls over the information systems that are the primary source for today's financial reporting requirements.

64.2.1 Section 404 of SOX. Since compliance with SOX is supported by the organization's ability to develop, implement, and manage a combination of business reporting requirements and management controls, many of the organization's information systems are directly or indirectly involved in the SOX-compliance process. The extent of involvement can range from end-user computing applications (e.g., spreadsheets) all the way to the data center. Business leaders must address these issues and must provide adequate support to achieve continuous compliance with SOX.

SOX addresses a number of areas requiring senior executive leadership and support. However, Section 404 of SOX is most relevant to information-security programs, and requires businesses to:

- Evaluate the adequacy of internal controls that impact financial reporting
- Implement new controls as necessary
- Annually test and report on the assessment results of internal controls

Section 404 requires the business to implement a comprehensive internal control framework sustainable throughout the enterprise. Furthermore, these controls must not only directly protect the integrity of financial data and indirectly protect the systems that support the data, but the business is also required to show that these controls are operational and are functioning as intended. Since compliance with SOX affects the organization at an enterprise-wide level, information-security protections should be designed, implemented, and managed from a holistic perspective supported by management as part of an overall business strategy, rather than as a single compliance activity.

In most organizations, almost all corporate information, including financial information, is linked to enterprise information systems. Many systems feed or share information with each other. Therefore, the organization must look at data origins, data flows, and data output to validate that they have adequately identified all of the information that could impact financial reporting. Management must provide the appropriate level of support for robust controls to protect the data resources of critical corporate information systems, which ultimately act as the primary data source the business must rely on when delivering and validating the integrity of the organization's financial statements.

With today's modern technologies pervasively deployed in business environments, organizations are likely to face a variety of SOX compliance challenges. These challenges result from the vast number of information systems, the associated supporting infrastructure, highly connected and communication intensive networks, and in many cases externally facing Web-based applications that are highly susceptible to exploitation.

Additional challenges come from inside as well. The insider threat is quickly becoming one of the most difficult security challenges to overcome, and causes the majority of

security problems for many organizations (see Chapters 13 and 45 in this *Handbook*). Internal data protection challenges are further complicated by the increase in *privileged users* such as database administrators, system administrators, and a variety of other users with powerful access to critical information systems. Such users often unknowingly create significant SOX compliance risks for their organization. These powerful privileges are often overused, and ultimately can be the pathway to noncompliance or malicious activity within the organization. Furthermore, many of these privileged users can manipulate system log files so that their malicious activity goes unnoticed or in some cases becomes untraceable (see Chapter 53).

64.2.2 Management Perspectives on SOX. Properly administered compliance initiatives must strike a balance between information security and system usability, to support business and compliance objectives without creating an environment of distrust. Although the insider threat is a significant threat to many businesses, the senior leadership team must accept that employees are a critical component of business operations. Since management controls are intended to protect sensitive system data and not to create a hostile work environment, SOX-related corporate policies and standards must support the spirit of the act by ensuring that organizations closely monitor and control access to critical business systems while not impeding employee productivity. Achieving this balance is not always an easy task.

Regardless of the specific strategies organizations employ to design, implement, and manage, there are some fundamental operational policies and processes that should be put in place to protect the integrity of both the data and ultimately the organization's financial reporting.

Fortunately, many organizations may find they already have a majority of the required controls in place, because many of the control requirements can be accomplished through traditional information security best practices. Some of the more common examples include controls that address segregation of duties (e.g., among development, test, and production environments) and appropriate data-security controls and safeguards.

A list of more common database safeguard components that senior management should focus on when complying with Section 404 of SOX follows.

- Monitor database access by privileged users
- Monitor changes in privileges
- Monitor access failures
- Monitor schema changes
- Monitor direct data access

In order to get a better grasp on all of the complexity associated with database compliance challenges, senior leadership teams in many organizations elect to take a back-to-basics approach. One of the most effective strategies to deploying such a simplified approach is to use the traditional five *W*s. Specifically, organizations should be asking these questions in terms of database compliance:

1. Who did something to the database?
2. What did they do?
3. Which specific data component?

4. When was the activity performed?
5. Where was the information accessed?

When members of the organization's senior leadership team approach the SOX compliance issue in this manner, they are in a better position to understand and ultimately to assume the responsibility for the data security and integrity for which they are accountable.

The requirement of having an organization's senior management work with an external auditor to report on the internal control program is, unfortunately, the most costly requirement of SOX and requires an enormous effort.

However, organizations should also take note that there are a number of benefits from designing, implementing, and managing effective compliance strategies, beyond initial SOX compliance. Establishing a comprehensive internal controls program not only aids in SOX compliance but also forms a foundation for operational improvements. Anytime the organization's senior management team can institute an effective internal controls program, it is not only taking steps to compliance, but also building a solid foundation on which enterprise-wide data governance programs can follow.

64.3 GRAMM-LEACH-BLILEY ACT. One of the most significant pieces of legislation to affect the financial services industry was the *Financial Services Modernization Act* of 1999, commonly referred to as the *Gramm-Leach-Bliley Act* or GLBA.[2] It was enacted in response to the standardization of the U.S. banking and insurance industries in the late 1990s. The financial services industry was undergoing substantial change during this period. Since financial institutions and insurance companies were allowed to merge and consolidate their operations, legislators and industry observers were concerned about providing adequate protection of consumer rights and data protection. Specifically, significant concerns developed over the consolidation of consumer data. There was substantial fear that sensitive consumer data would be openly shared among financial organizations and their subsidiaries. This open environment would likely threaten consumer rights and the security of sensitive and personal financial data. The consolidation and reform within the financial services industry was enabling the creation of new financial services holding companies that could offer a full range of financial products. Prior to the passage of GLBA, some major financial institutions were selling personal and sensitive detailed customer information to business partners. This type of disclosure often included the disclosure of account numbers and other highly sensitive data to telemarketing firms. These telemarketing agencies often used the account numbers to charge customers for products and services they did not want and that had no value to customers.

64.3.1 Applicability. GLBA applies to U.S. domestic financial institutions; it defines financial institutions, as "companies that offer financial products or services to individuals, like loans, financial or investment advice, or insurance." GLBA coverage includes but is not limited to organizations that provide insurance, securities, payment settlement services, check cashing services, credit counselors, and mortgages.

GLBA is intended to address the proper handling of nonpublic financial information. However, senior management teams of financial institutions must acknowledge that the act also includes a wide range of information that is not obviously financial in nature.

This additional coverage is intended to offer further protections to the consumer and properly align with the spirit of the law. The additional types of information that must be protected include the consumer's name and address. However, protection may extend to:

- Information given to a financial institution in order to receive a financial product or service
- Information generated or remaining as a result of a transaction between a financial institution and a consumer
- Information obtained by the financial institution while providing a financial product or service to a consumer

GLBA requires financial institutions to safeguard nonpublic personal information (such as a Social Security number, credit card details, or a bank account number) provided by a consumer under various privacy rules or resulting from a transaction or other service performed on behalf of the consumer. This information is not necessarily considered financial information. For example, under GLBA, a consumer's name identifying a recipient of services from a specific institution is also considered nonpublic information that must be protected under GLBA. Specifically, GLBA mandates:

- Secure storage of consumer personal information
- Providing adequate and sufficient notice to consumers regarding how the financial institution shares their consumer personal and financial information
- Providing consumers with the choice to opt out of sharing their personal and financial information

64.3.2 Enforcement. GLBA is enforced by the Federal Trade Commission (FTC), various financial industry regulators, and state attorneys general. More stringent requirements may be placed on the financial institutions by individual states, because GLBA does not preempt state law.

Although current legislation does not offer a remedy of civil action, a financial institution's failure to comply with notice is considered a deceptive trade practice by both state and federal authorities. Some states do have specific legislation that offers a private right of action for consumers.

Title V of the act provides that financial institutions may share practically any information with affiliated companies but may share information with nonaffiliated companies for marketing purposes only after providing an opportunity for the consumer to *opt out* of the information sharing process. Management needs to play a proactive role in providing administrative, managerial, and technical support to achieve compliance with Title V of the Act. Consumers must be provided with easy-to-understand and easy-to-use opt-out choices. Organizations that make opting out difficult, or that otherwise appear to be duplicitous, are likely to be viewed as noncompliant with Title V of GLBA.

The provisions of Title V do not preempt or supersede state law. If organizations are conducting business in a state that has more stringent requirements, then those state requirements must be met, in addition to GLBA. Section 505 requires that the act and its associated regulations be enforced by various federal and state regulatory agencies having jurisdiction over financial institutions. This requirement of GLBA

offers federal enforcement authority to a variety of enforcement agencies. The more common enforcement agencies include:

- Office of the Comptroller of the Currency
- Federal Reserve Board
- Federal Deposit Insurance Corporation
- Securities and Exchange Commission

An organization's senior leadership team should also note that the FTC has general enforcement authority for any financial institution that does not fall within the jurisdiction of any of the specific enumerated regulatory agencies.

As previously stated, the intent of GLBA is to offer adequate protection to consumers conducting business with financial services firms. Although the hope is that organizations will comply with GLBA on goodwill, the act was created with some serious repercussions to financial institutions that do not comply. GLBA does not include a private right of action, but financial institutions are required to give consumer notice, and they could face liability under deceptive trade practice statutes if the notices are determined to be deceptive or inaccurate. Furthermore, financial institutions that fail to comply with GLBA may also be subject to penalties under the *Financial Institution Reform Recovery and Enforcement Act* (FIRREA). FIRREA contains and offers penalties that range from up to $5,500 for violations of laws and regulations; up to $27,500 if violations are unsafe, unsound, or reckless; and up to $1.1 million for knowing violations. Given the consequences of noncompliance, organizations need to take the appropriate actions to obtain consumer protection and compliance with GLBA.

64.3.3 Consumers and Customers. One common misconception of GLBA is the interpretation of the legislation as it relates to the nonconsumer customers of financial institutions. This is likely to occur since the terms *consumer* and *customer* are often used interchangeably. However, in the context of GLBA, the critical distinction must be made. GLBA defines a *consumer* as "an individual who obtains, from a financial institution, financial products or services which are to be used primarily for personal, family, or household purposes, and also means the legal representative of such an individual." A *customer* is a type of consumer who has established some sort of ongoing relationship with an institution accountable to GLBA. An example of a customer would be an individual who has received credit financing to purchase a car. This individual would need to make monthly payments to a financial institution and hence has established an ongoing relationship. Under this definition, a business could not be defined either as a customer or as a consumer (since a business is not an individual[3]) and therefore does not fall under the protections of GLBA.

64.3.4 Compliance. Compliance with GLBA signifies that financial institutions must comply with a number of the act's provisions. From a high-level perspective, organizations must adhere to these points when seeking compliance with GLBA:

- Organizations must provide consumers with clear and conspicuous notice of the financial institution's information-sharing policies and practices.
- This notice must be given at the start of a new customer relationship and maintained on an annual basis.

- Organizations must provide customers with the right to opt out of having their nonpublic personal information shared with nonaffiliated third parties (unless the activity falls under one of the GLBA exceptions).
- They must refrain from disclosing to any nonaffiliated third-party marketer, other than a consumer reporting agency, an account number or similar form of access code to a consumer's credit card, deposit, or transaction account.
- Financial institutions must comply with the regulatory standards established to protect the security and confidentiality of customer records.
- Financial institutions must also protect against security threats and unauthorized access to such protected customer information.

64.3.5 Privacy Notices. As a part of GLBA, financial institutions are required to disseminate privacy notices to their consumers explaining what information the financial institution will collect about the consumer, whom the information may be shared with, and how the financial institution protects that information. The information the notice must refer to is the consumer's nonpublic personal information (such as a Social Security number, bank account number, or credit card number) collected, for example, during an application for auto financing. In an effort to help companies comply with this requirement in a consistent and effective manner, the FTC (in conjunction with various other financial regulators) established a GLBA privacy notice standard. Under this standard, financial institutions must provide a privacy notice at the inception of a relationship with a consumer and once a year for as long as the relationship persists. Furthermore, this notice must provide consumers with the option to declare within 30 days after the receipt of the notice that they do not want their information shared with the third parties mentioned in that privacy notice. Financial institutions must also mandate that the privacy notice itself be written in such a way that it is a clear and accurate statement of the company's privacy practices. Provided these requirements are met by the financial institution's privacy notice, the institution may share consumer information with its affiliated companies.

A financial institution may also share consumer information with nonaffiliated organizations provided it has clearly disclosed what information will be provided, to whom, and how it will be protected, and provided that consumers have a method of opting out clearly described in their privacy notices. However, GLBA prohibits financial institutions from disclosing consumer account numbers to nonaffiliated companies for purposes of telemarketing, direct-mail marketing (including through email), even if the consumer has not opted out of sharing the information for marketing purposes. There will be special circumstances in which consumers will not have the option to require that their information not be shared. Some of the more common occurrences of these special circumstances are:

- When a financial institution is required to share information with an outside organization as part of fulfilling its customer obligations (e.g., data processing services)
- When a financial institution is legally required to share the information
- When the information is shared with an outside service provider that market the products or services of the financial institution

64.3.6 GLBA Safeguards Rule. A significant component of GLBA is the *GLBA Safeguards Rule*. The Safeguards Rule is intended to support the privacy and

protective rules within GLBA. In many organizations, it is the operational driver supporting the other components of GLBA. If the other components of GLBA are the *why* of the legislation, then the Safeguards Rule is the *how*.

GLBA contains specific requirements and safeguards that are applicable to both electronic and paper records. Management must protect records in both formats, even though the widespread use of electronic data processing makes it easy to overlook sensitive consumer information in hardcopy format because an increasing number of employees work almost exclusively with electronic representations of the data. GLBA requires financial institutions to develop, implement, and continuously manage a *comprehensive* information security program, which it defines as including administrative, technical, and physical safeguards to protect the security of customer information.

64.3.7 Flexibility. GLBA covers a broad array of financial institutions, and therefore mandates that an information security program be appropriate to the size, complexity, nature, and scope of the activities of each specific financial institution. This level of flexibility offers an organization's management team the flexibility to balance consumer protection against business objectives. However, it is critical that this flexibility not be abused and that the intent of the act be addressed. Furthermore, an organization's senior leadership team must not interpret the lack of specific safeguards as an opportunity for noncompliance with GLBA, but rather as an opportunity to provide consumer protection without impeding business operations.

Flexibility aside, certain key requirements must be followed to be compliant with the Safeguards Rule. The foundational security practices include having at least one designated employee to audit systems; determine risks; and develop, implement, and manage procedures to address information security. Specifically, a financial institution must provide three levels of security:

1. **Administrative security.** Includes program management of workforce risks, employee training, and vendor oversight
2. **Technical security.** Includes technical controls for computer systems, networks, applications, access controls, and encryption
3. **Physical security.** Includes safeguarding facilities, corresponding environmental protection, and disaster recovery protections

In order for financial institutions to be fully compliant with the Safeguards Rule, they need to create comprehensive internal controls for strong administrative, technical, and physical safeguards. To help guide an organization's senior leadership team toward GLBA compliance, some of the major requirements that financial institutions must keep in mind when building an organizational culture founded on consumer protection and GLBA compliance are listed next:

- Validate that consumer information is kept secure and confidential.
- Validate that consumer information is protected from likely threats to its security and integrity.
- Validate that consumer information cannot be accessed by any unauthorized entities, or accessed in a way that would result in substantial loss or in a way that would inconvenience the consumer.[4]

EXHIBIT 64.2 Recommended Evaluation Procedures

Key Questions or Considerations	
I.	**Determine the involvement of the board.**
A.	Has the board or its designated committee approved a written Corporate Information Security Program that meets the requirements of the Information Security Guidelines?
B.	If the board has assigned responsibility for program implementation and review of management reports to an individual or to a committee, is the necessary knowledge, expertise, and authority to perform the task available?
C.	Does the program contain the required elements? 1. If more than one information security program exists for the institution, are the programs coordinated across organizational units?
D.	Are the reports from management to the board (or its designated committee) useful? Does the report adequately describe the overall status of the program, including material risk issues, risk assessment, risk management, and control decisions, service provider oversight, results of testing, security breaches and management's response, and recommendations for program changes? 1. How often does the board (or its designated committee) review reports?
E.	Overall, do management and the board (or its designated committee) adequately oversee the institution's information security program?
II.	**Evaluate the risk assessment process.**
A.	Review the risk assessment program. 1. How does the institution assess risk to its customer information systems and nonpublic customer information? 2. Has the institution evaluated the risk to the entire customer information system? 3. Has the institution used personnel with sufficient expertise to assess the risk to its systems and customer information on an enterprise-wide basis? 4. Is the risk assessment part of a formal risk assessment process with timelines and milestones? If not, how will management ensure timely completion? 5. Does the institution have a process for identifying and ranking its information assets (data and system components) according to sensitivity? How does it use this process in its risk assessment?
B.	Assess adequacy and effectiveness of the risk assessment process. 1. Does the institution identify all reasonably foreseeable internal and external threats that could result in unauthorized disclosure, misuse, alternation, or destruction of customer information, or of customer information systems? 2. Does the institution support its estimate of the potential damage posed by various threats? 3. Review the institution's existing controls to mitigate risks. Does the institution's analysis consider the current administrative, physical, and technical safeguards that prevent or mitigate potential damage? 4. Does the institution use test results to support its assessment of the adequacy and effectiveness of those controls?

(*continued*)

EXHIBIT 64.2 (*Continued*)

Key Questions or Considerations	
C.	Does the institution identify and prioritize its risk exposure, decide on the risks it must mitigate, and create a mitigation strategy? Is the decision to accept risks documented and reported to the appropriate management levels? 1. Does the institution promptly act to mitigate risks that pose the immediate possibility of material loss? 2. How does the institution demonstrate that the mitigation strategy was reviewed by appropriate officials? 3. Does the risk assessment provide guidance for the nature and extent of testing? 4. Does the risk assessment include vendor oversight requirements?
III.	**Evaluate the adequacy of the program to manage and control risk.**
A.	Review internal controls and policies. Has the institution documented or otherwise demonstrated, at a minimum, that it considered the following controls, and adopted those it considered appropriate? 1. Access controls, such as controls to authenticate and permit access to customer information systems to authorized persons only. 2. Access restrictions at physical locations, such as buildings and computer facilities, to permit access to authorized persons only. 3. Encryption of electronically transmitted and stored customer data. 4. Procedures to ensure that systems modifications are consistent with the approved security program. 5. Dual control procedures, segregation of duties, and employee background checks. 6. Monitoring systems and procedures to detect actual and attempted attacks on, or intrusions into, customer information systems. 7. Response programs specifying actions to be taken by specific individuals when the institution suspects unauthorized access (i.e., incident response). 8. Measures to protect against destruction, loss, or damage of information from potential environmental hazards, such as fire and water damage or technological failures.
B.	Is staff adequately trained to implement the security program? 1. Obtain from management a listing of the training provided to all users of the institution's system.
C.	Determine whether key controls, systems, and procedures of the information security program are regularly tested by independent third parties, or by qualified independent staff, in accordance with the risk assessment. 1. Assess whether the nature and frequency of testing is consistent with the risk assessment. 2. Assess whether tests are conducted or reviewed by independent third parties or qualified staff independent of those that develop or maintain the security program. 3. Assess whether management reviews test results promptly. Assess whether management takes appropriate steps to address adverse test results.

EXHIBIT 64.2 (*Continued*)

Key Questions or Considerations	
IV.	**Assess the measures taken to oversee service providers.**
A.	Determine whether the institution exercises due diligence in selecting service providers.
B.	Determine what information is supplied to service providers.
C.	Obtain a copy of the contract(s) with the service provider(s). Determine whether contracts require service providers to implement appropriate measures to meet the objectives of the guidelines.
D.	If the institution's risk assessment requires monitoring a service provider, then perform the following steps for each applicable service provider. 1. Determine whether the service provider contract provides for sufficient reporting from the service provider to allow the institution to evaluate appropriately the service provider's performance and security, both in ongoing operations, and when malicious activity is suspected or known. 2. Determine whether the institution's actions adequately control information supplied to service providers, ensuring that the information is managed and secured properly. 3. Review financial condition of service provider.
V.	**Determine whether an effective process exists to adjust the information security program.**
A.	Does the institution have an effective process to adjust the information security program as needed? Is the appropriate person assigned responsibility for adjusting the information security program?
B.	Review procedures that are in place to ensure that when the institution makes changes in technology, and in its business functions, the requirements of the guidelines are also considered. These changes can include: 1. Technology changes (e.g., software patches, new attack technologies, and methodologies). 2. Sensitivity of information. 3. Threats (both as to nature and extent). 4. Upcoming changes to the institution's business arrangements (e.g., mergers and acquisitions, alliances and joint ventures, outsourcing arrangements). 5. Upcoming changes to customer information systems (e.g., new configurations or connectivity, new software).
C.	Determine whether appropriate expertise is applied to evaluate whether changes to the information security program are necessary.
D.	Determine whether appropriate controls exist to ensure that changes to the information security program are properly implemented in a timely, risk-minimized manner.
VI.	**Summarize and communicate your findings.**
A.	Discuss issues, conclusions, and potential violations with executive in charge.
B.	Discuss findings with institution management. If you have identified material issues, obtain and document management commitments to address those Issues.
C.	Complete work papers.
D.	Detail findings with support in a Summary Comment.

Fortunately, for many organizations, implementing the Safeguards Rule does not usually require significant structural changes to their information security program. This occurs because the principle of the Safeguards Rule requires compliance with basic information-security program elements. These elements are based on best practices that should already be a part of an effective information-security program.

Specifically, each financial institution must designate an employee to manage the programs safeguards. Depending on financial institution's size, culture, and organizational structure, the actual title and responsibilities of the position will vary significantly. However, at a high level, this function typically requires the identification and assessment of the organization's risk as it relates to consumer information. Furthermore, the organization's risk management must be carried out diligently, so that each relevant area of the institution's operation successfully completes a comprehensive evaluation of the effectiveness of its safeguards. In most cases, the effectiveness of the organization's safeguards can be measured by the organization's ability to manage risk successfully by designing, implementing, and managing adequate information security controls. This activity should further be supported by senior management's ability successfully to monitor and test the information risk-management program on a regular basis.

Finally, it is important that the financial institution's senior leadership team proactively protect customer information, even when working with third parties. Under GLBA, financial institutions must also select the appropriate service providers to support their GLBA requirements and to enter into security-minded agreements. Simply utilizing third parties does not free financial institutions from their GLBA obligations. The organization's senior leadership team must establish agreements in a manner that adequately protects consumer information and that requires the service provider to implement appropriate and GLBA-relevant safeguards. GLBA stresses that, at a minimum, the safeguards performed by the third party must be equal to the protections offered internally by the financial institution's senior leadership team.

64.4 EXAMINATION PROCEDURES TO EVALUATE COMPLIANCE WITH GUIDELINES FOR SAFEGUARDING CUSTOMER INFORMATION.

The Office of the Controller of the Currency of the United States Department of the Treasury has published a useful table of recommended evaluation procedures that will help readers apply the principles discussed in this chapter.[5] A simplified and edited representation of the procedures is supplied in Exhibit 64.2.

64.5 CONCLUDING REMARKS.

An overall increase in regulatory requirements is likely to emerge in coming years. Some of these requirements may be industry specific or broad enough to reach all organizations. Good security professionals will build and maintain comprehensive governance, risk, and control programs. This will give their organizations a strong foundation for any future regulatory requirements.

64.6 FURTHER READING

Anand, S. *Essentials of Sarbanes-Oxley*. Wiley, 2007.

Aspatore Publishers, eds. *Recent Trends in Privacy and Data Security: Leading Lawyers on Analyzing Information Storage Regulations and Developing Effective Data Protection Policies.* Thomson Reuters Westlaw, 2013.

Bianco, K. M., J. Hamilton, K. R. Benson, A. A. Turner, and J. M. Pachkowski. *Financial Services Modernization: Gramm-Leach-Bliley Act of 1999 Law and Explanation*. CCH Inc., 1999.

Buchanan, Jim "Cloud Computing: 4 Tips for Regulatory Compliance." *CIO.com* 2011. www.cio.com/article/687434/Cloud_Computing_4_Tips_for_Regulatory_Compliance

Cohen, H. R., and W. J. Sweet. *After the Gramm-Leach-Bliley Act: A Road Map for Banks, Securities Firms, and Investment Managers*. Practising Law Institute, 2000.

Cox, F. D. *Information Security: Risk Management of GLBA Privacy and Service Provider Oversight*. Amazon Digital Services, 2011.

Crosman, Penny. "Compliance Roadmap for Proposed Social Media Rules for Banks" *American Banker* (2013). www.americanbanker.com/issues/178_102/compliance-roadmap-for-proposed-social-media-rules-for-banks-1059418-1.html

Dunham, W. B. *After the Gramm-Leach-Bliley Act: A Road Map for Insurance Companies*. Practising Law Institute, 2000.

Greene, E. F., D. M. Becker, and L. N. Silverman. *Sarbanes-Oxley Act: Analysis and Practice*. Aspen Law & Business, 2003.

Hermann, D. S. *Complete Guide to Security and Privacy Metrics: Measuring Regulatory Compliance, Operational Resilience, and ROI*. CRC Press, 2006.

Kearney, E. F., R. Fernandez, J. W. Green, and D. M. Zavada. *Wiley Federal Government Auditing: Laws, Regulations, Standards, Practices, and Sarbanes-Oxley*, 2nd ed. Wiley, 2013.

Lekatis, G. *Understanding Sarbanes-Oxley: What Is Different after July 2013*. Amazon Digital Services, 2013.

Miles, B. L. *The Canadian Financial System*. Nova Science Publishers, 2004.

Moeller, R. R. *Executive's Guide to COSO Internal Controls: Understanding and Implementing the New Framework.* Wiley, 2013.

Tarantino, A. *Manager's Guide to Compliance: Sarbanes-Oxley, COSO, ERM, COBIT, IFRS, BASEL II, OMB's A-123, ASX 10, OECD Principles, Turnbull Guidance, Best Practices, and Case Studies*. Wiley, 2006.

White House. "Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy." White House Website, February 2012. www.whitehouse.gov/sites/default/files/privacy-final.pdf

64.7 NOTES

1. American Institute of Certified Public Accountants Center for Audit Quality, www.aicpa.org/INTERESTAREAS/CENTERFORAUDITQUALITY/RESOURCES/Pages/Resources.aspx; and see "Section 404(b) of Sarbanes-Oxley Act of 2002," www.aicpa.org/Advocacy/Issues/Pages/Section404bofSOX.aspx
2. Gramm-Leach-Bliley Act: 15 USC SubChapter 1, §§6801–6809, www.ftc.gov/privacy/glbact/glbsub1.htm
3. At least, not in this context. For a different perspective on the question of organizations as people, see discussions of Citizens United v. Federal Election Commission, www.scotusblog.com/case-files/cases/citizens-united-v-federal-election-commission

4. FTC, "Standards for Safeguarding Customer Information: Final Rule," May 23, 2002, 16 CFR §314 ¶1a, www.ftc.gov/os/2002/05/67fr36585.pdf
5. U.S. Department of the Treasury, Office of the Controller of the Currency, "Examination Procedures to Evaluate Compliance with the Guidelines to Safeguard Customer Information," OCC- 2001-35 (2001-07-18). www.occ.gov/news-issuances/bulletins/2001/bulletin-2001-35.html

CHAPTER 65

THE ROLE OF THE CISO

Karen F. Worstell

65.1 CISO AS CHANGE AGENT 65·1

65.2 CISO AS STRATEGIST 65·3

65.2.1 Reliance on Digital Information 65·4
65.2.2 Inherent Insecurity of Systems 65·5
65.2.3 World Trends 65·5

65.3 STRATEGY, GOVERNANCE, AND THE STANDARD OF CARE 65·6

65.3.1 Standard of Care 65·6
65.3.2 Governance and Accountability 65·9
65.3.3 Roles and Responsibilities 65·11
65.3.4 Reporting 65·13
65.3.5 Monitoring 65·13
65.3.6 Metrics 65·13
65.3.7 Executive Visibility 65·13

65.4 SUMMARY OF ACTIONS 65·13

65.5 RECOMMENDATIONS FOR SUCCESS FOR CISOs 65·14

65.5.1 Education and Experience 65·15
65.5.2 Culture of Security in the Business 65·15
65.5.3 Alliance with Corporate and Outside Counsel 65·16
65.5.4 Partnership with Internal Audit 65·16
65.5.5 Tension with IT 65·17
65.5.6 Organizational Structure 65·17
65.5.7 Responsibilities and Opportunities outside of CISO Internal Responsibilities 65·18

65.6 CONCLUDING REMARKS 65·18

65.7 FURTHER READING 65·19

65.8 NOTES 65·19

65.1 CISO AS CHANGE AGENT. The title of chief information security officer (CISO) has evolved because of the realization that the function of the chief information officer (CIO) is so broad as to require another person to focus specifically on the *security* elements of information. Another motivation derives from the fact that the CISO can perform functions that are not usually associated with the CIO. Our approach to information security needs to change in response to the disruptive events affecting the network and the boardroom. CISOs should be the change agents to make this happen. This is a shift from the majority of CISOs' emphasis today as senior managers of information technology (IT) security.

Today, CISOs are in the trust business, due to the need to create and maintain a network of trust among all the people, business processes, and technology of an enterprise and its partners. The interconnected ecosystem that developed since the commercialization of the Internet has seen dramatic shifts of trust: Consumers are thinking twice

before conducting business transactions online, and governments have intervened with regulations to improve the trust environment. The shift has between toward less trust each time we encountered a *disruptive event*—an event that awakened us to the fact that things were not as they seemed and made us aware that protection of networked information and systems needed to address something we had not anticipated. The timeline of disruptive security events that began in 1986 with the promise of productivity by Macs and PCs worsened with the Brain virus (1986). Trust continued to decline with the Morris Worm (1988) and the Concept.A macro virus (1995). Our world became flatter,[1] and security became an entirely new kind of problem with the commercialization of the Internet and the introduction of the Web. In 1995, we still trusted the networked ecosystem. We saw the malicious works of nefarious individuals and groups, and we saw the future threat, but still did not fully realize how bad it would get.

Throughout this chapter, the author refers to *information security* and the *CISO*. This is not to the exclusion of aspects of security that do not deal directly with information systems; rather, it is intended in the broadest possible sense to address the interdependent disciplines that are required successfully to protect information in all its forms. It is in this context that the title *CISO* is used, inclusive of *CI* and *CSO* rather than exclusive.

In 2001, the horrific events of 9/11 rocked the planet, Enron collapsed, trust plummeted with Slammer (2002), and the insidious stealthy vectors that utilize port 80 to steal everything from personal identities to bank accounts and company data. The networked ecosystem has provided no end of new material for our profession while making *computer security*, *spyware*, and *identity theft* household terms.

Mark Twain once said the definition of insanity is to do the same thing over and over and expect different results. It is clear that we will need to apply a different way of thinking and new roles and responsibilities to the discipline of information security if we are to make progress in enabling the kind of protection that is required by businesses and agencies while enabling business agility and competitiveness. This is the opportunity of the CISO.

Trends reveal that we are just past the cusp of a new kind of *disruptive event* worldwide—the introduction of legislative and regulatory mandates to ensure effectiveness of controls for protection of consumers, critical infrastructures, and shareholder value. We deal today with a myriad of high-impact network-based crimes: extortion, corporate espionage, and massive fraud. Many experienced CISOs believe that a *cyber–Pearl Harbor* scenario is plausible. Since 2001, we are acutely aware of the vulnerability of critical infrastructures, at least 85 percent of which are owned and maintained by the private sector. Recent figures[2] indicate that new home PCs are compromised with anything from spyware to Trojans within 5 minutes of connecting to the Internet (it takes up to 60 minutes to download protective software). Although corporate PCs may enjoy a relative degree of protection, as many as 90 percent of home PCs have some kind of spyware. Corporations with the resources to address information security have doubled the use of IT security standards and guides since 2003. But the insidious attack on the home user is winning a game with serious stakes: Consumer confidence in the Internet is waning. Once *computer security* was the domain of techies and *geeks*; it is now pervasive enough to be a household concern.

Security has definitely reached the awareness of the C-level offices (chief executive officer, chief financial officer, etc.) on mahogany row, and there are articles in the popular press on a daily basis concerning information security around the world. This visibility ensures that no one can hide behind a plea of ignorance of security problems.

Legislators and regulators are more concerned than ever, and new statutes, codes, regulations, and guidelines have proliferated. In the face of huge threats and the need for a corporate duty of care, network security remains a concern, but compliance to regulations and statutes and the ability to demonstrate adherence have become mandatory.

Penalties for noncompliance to these recently established governing rules are setting breathtaking precedents: $1.2 billion for a large U.S. financial institution, $15 million for another. The U.S. Veteran's Administration admitted to the compromise of personal information, including Social Security numbers, of 26.5 million living American veterans.[3] It may be noted that even New Zealand was for sale on eBay® recently—reportedly the efforts of an unknown party in Queensland—and was taken off line after 22 bids had reached A$2,000.[4] The reserve, if any, was probably not met. We live in a new world. Trust is no longer assumed; it is easily broken in the most inadvertent to the most imaginative ways. The case law for settlements and summary judgments in information security matters is providing the basis for the best potential dollar-based impact analysis on business and information-security risk since Basel II, which is the second set of international guidelines on banking laws and regulations recommended by the Basel Committee on Banking Supervision.[5]

Companies, and clearly U.S. Government agencies, do not have control of intellectual property, employee data, or consumer data. The need to retain corporate records in the event of litigation regarding business matters, and to produce such information on demand under discovery, is infeasible for many firms with petabytes or more of structured and unstructured data. Recent legislation, regulations, and case law will drive the need for information security programs to solve this enormous data management and security problem and to demonstrate a level of effectiveness to a standard of care that will stand up to the scrutiny of opposing counsel. It is the role of the CISO to understand fully the implications of these new burdens and to incorporate them into a strategy along with the business strategy of each enterprise. Where information security spending was benchmarked by Gartner Group at 1 to 3 percent of overall IT budget in the mainframe days, and 5 to 7 percent at the dawn of distributed computing, the spending to resolve problems created by past failures to address solid controls and to be prepared for the threats of the future to any degree should well exceed 10 percent of an overall IT budget. This does not include the special allowances that will be required to treat some of the worst offenses. According to the 2004 Information Security survey conducted by PricewaterhouseCoopers and *CIO* magazine, the best-practices group spent 14 percent of their IT budgets on information security each year.[6] Companies must plan for these expenditures; the role of the CISO will be to help clarify obligations, business necessities, and strategic spending.

65.2 CISO AS STRATEGIST. Information is the prize that motivates wrongdoers, perpetrators, and miscreants who tamper with, destroy, and penetrate systems that process, store, and transmit digital information. Phil Condit, former chairman of The Boeing Company, said it well at his keynote to the International Information Integrity Institute in 1995: "Information is the business."[7] His statement was visionary at the time. Indeed, since 1995, it is safe to say that information is the foundation of business, but it is the way it is used, combined, mined, shared, represented, accessed, and manipulated for business logic's sake that creates competitive advantage. And each of these areas of application logic and information requires rules-based protection to ensure that the business is basing decisions on sound knowledge.

There are three main drivers for CISOs to define a new strategy for information security:

1. Systems are inherently insecure as a result of enormous variations in configuration, sheer complexity, and volume of vulnerabilities. A risk-based approach using conventional methods is no longer valid (if it ever was).
2. The reach of global business processes, personnel, and business systems introduces new considerations into an already complex security problem.
3. The playing field between the *protectors* and the *interlopers* is dramatically uneven. There are literally hundreds of thousands, if not millions, of interlopers with a follow-the-sun 7×24×365 factory of attacks. Companies cannot afford any such investment to counter the attacks, even if it were possible with a combination of people and automation. The attack vectors change too quickly to be able to get ahead of the curve on the protection side of the equation. Nation-states, organized crime, terrorist organizations, fraudsters, and identity thieves are continuously developing ingenious models to attack, steal, and destroy precious assets of individuals as well as of enterprises.[8] Worse, professionals continue to consider trusted insiders as the most significant threat. With a globally distributed workforce of contractors, employees, and suppliers, it is necessary to have an approach to information protection controls that enables a reasonable assurance that both the external and insider threats are appropriately addressed.

The CISO as a strategist will be successful adopting and integrating new methods into the business, such as a rules-based standard of care and due diligence to that standard of care. This will necessarily include existing methods of network protection, data classification, and so forth—the standard of ISO/IEC 17799:2005 provides an excellent framework. It is the decision-making process about priorities that must change.[9]

As a strategist, the CISO will need to look at the security problem as an executive businessperson would. Adopting the kinds of strategic thinking described in classics such as *The Art of the Long View* is a key success factor.[10] A partial list of things to consider includes the reliance on information, why protection is important, the insecurity of systems, the futility of risk-based security, and world trends.

65.2.1 Reliance on Digital Information. Digital information is the lifeblood of our commerce, financial infrastructure, healthcare, transportation, energy, and even our very identities as citizens. In the United States alone there are 15 distinct critical infrastructures, vital to the interests of the country and to national security:

1. Information technology
2. Telecommunications
3. Chemicals
4. Transportation systems
5. Emergency services
6. Postal and shipping services
7. Agriculture and food
8. Public health and healthcare
9. Drinking water/water treatment

10. Energy
11. Banking and finance
12. National monuments and icons
13. Defense industrial base
14. Key industry/technology sites
15. Large gathering sites[11]

Eight-five percent of the systems that make up these critical infrastructures are owned by private enterprise. The systems that interconnect to deliver services require integrity, availability, the trusted relationships, and confidentiality. If we cannot trust the systems, we cannot trust the information. If we cannot trust the information, we cannot trust decisions based on that information. Information security also now requires us to be able to demonstrate digital ownership and even chain of possession. We have to maintain digital information in such a way that we know where our IP is at all times, what records our company must maintain, and how to store, label, and retrieve them. The definition of information security as CIA—confidentiality, integrity, and availability—is too simplistic.

65.2.2 Inherent Insecurity of Systems. Systems include hardware, software, utilities, scripts, and transport media, all of which are ultimately created by humans with scheduled deadlines and constrained budgets, and installed by organizations with scheduled deadlines and constrained budgets. It is all flawed. The vulnerabilities that exist today, let alone the ones we are going to learn about tomorrow and the next day, cannot all be addressed. Perfect security is a myth, unattainable, and arguably a waste of time to try. Risk-based methods for identifying required security tactics do not scale, nor is there meaningful data that describe probability and ALE (annualized loss expectancy) realistically for information security.

65.2.3 World Trends. Our world is changing. As summarized in a report titled "Ten Trends to Watch in 2006" published by McKinsey and Company,[12] trends that have already had significant implications for security professionals include:

- Dramatic geographic shifts in centers of economic activity, particularly in IT services, where labor and talent are increasing globally. Worldwide, distributed, and devolved security models will be required, as the supply chain and internal processes in a company's value chain are widely distributed. The implications of physical security, information systems protection, personnel practices, and the diversity of governing regulations and statutes require a much broader, business-based view of asset protection. Use of company-owned, leased, and outsourced facilities will prove to be a challenge to asset protection.
- Technological connectivity will transform the way we live and interact and will completely disrupt current security infrastructures. Enhanced connectivity and mobility using smaller (i.e., less observable), highly capable Internet-enabled devices requires access to enterprise information resources from anywhere, at any time, putting an enormous strain on rule sets designed to filter incoming and outgoing information, and where it travels. Peer-to-peer networking, tiny mass storage devices, blending of personal and enterprise computing on common devices—these are just a few of the near-term changes that will completely alter

the way enterprise rules are handled for information protection. Grid computing, virtual machines, work from home—all are issues that challenge the notion of comprehensive asset protection, both physical and logical.

- New models of in-the-cloud knowledge, including production, access distribution, and ownership are emerging, and fundamental trust models remain to be defined or even described. *Software as a service* (SAAS) is already enabled for consumers and small businesses, and, probably, enterprise services will follow. The download and introduction of software over the Web, and access to proprietary information in an anywhere-computing environment, requires redefinition of the way we think about intellectual property and data management.

A coordinated, interdisciplinary management approach to a proper set of controls based on business risk is essential to deal effectively with the enormous requirements to protect information in business today. Whether the enterprise is private or public, regardless of nationality, a solid information protection strategy must recognize the implications of that principle.

To summarize, successful use and management of information is the business. Setting priorities that demonstrate due diligence to a properly business-driven standard of care for the confidentiality, integrity, availability, ownership, and proper possession of information is the charge of the CISO. It is essential that the CISO function at the level of executive management and as a business strategist, participating with the executive leadership team to enable the integration of due diligence to a standard of care into all business streams.

65.3 STRATEGY, GOVERNANCE, AND THE STANDARD OF CARE. Recognizing the information security changes that drive the role of the CISO, what should be done to be successful? What are the key focus areas that should define the role of a CISO?

Success for information security professionals will depend on several key factors in this new world:

- Standard of care (e.g., rules-based) strategy
- Governance and accountability
- Clear roles and responsibilities
- Metrics, reporting, and executive visibility

The remainder of this section examines each of these more fully. For a general discussion of management's role in information assurance, see Chapter 63 in this *Handbook*.

65.3.1 Standard of Care. A CISO has one vision statement that drives strategy: to establish due diligence to a standard of care for the business; that is to say, to put in place the mechanisms (controls, oversight, monitoring, metrics, and reporting) that will enable the business to demonstrate due diligence to that standard of care. Instead such a standard of care is rules based, but it is not prescriptive in itself. The standard of care is the set of documented business risks from an information protection perspective, and the declaration of a set of rules in policy, that effectively mitigate those business risks. The standard of care declares *what* will be done. It is the governing book of policy to be reviewed and monitored by the executive leadership team, the

audit committee of the board of directors, and senior management. It is the set of rules with which all internal standards and desktop procedures must comply.

Although a complete treatment of the identification of business risk is outside the scope of this chapter, it can be summarized in this way: Prior to identifying the controls within a standard of care, steps must be taken to qualify business risk relatively as *high*, *medium*, or *low* based on magnitude of potential business impact and perceived exposure. It does not use quantitative risk analysis methods such as those found in Octave, FAIR, or any other popular risk-based quantitative methodology. It does not use ALE. It does enable the business to defend the design of the standard of care, the controls of which are derived from the business risk analysis.

The standard of care derives its meaning from two primary sources: internationally accepted standards and the risk analysis that defines the business risks.

The basic set of internationally accepted standards is summarized in Exhibit 65.1. For more detailed analysis of security standards, see Chapters 44 and 51 in this *Handbook*.

It is important for these standards to be used as a foundation for the policy, to ensure that security policy is recognized by the larger body of security professionals, and to serve as an interchangeable trading partner agreement in the communication and enforcement of security expectations.

It is not in the scope of this chapter to go into the risk assessment methodology in detail; however, a general description will serve to indicate that this is not a vulnerability assessment. The risk assessment process is one used not to identify system vulnerabilities but to identify potential areas relating to people, process, and technology that could result in an exposure that reaches a defined threshold of business impact. For example, in a technology firm, one might align three major areas of concern: major inaccuracies in financial reporting (the domain of the Sarbanes-Oxley Act of 2002); failure to control IP adequately and the subsequent risk to trade secrets, patents, or copyrights; and failure to protect customer data adequately. Each of these slices of business risk can be the foundation for the questions in the risk assessment to ascertain major systems exposure, business process exposure, or exposure caused by people doing things incorrectly, such as human error or even malfeasance. Risk rankings will need to be on the basis of assumptions and figures acceptable to the senior leadership team; the wildcard in any risk ranking is probability. The business impact threshold will vary widely with each business. Thresholds may be set using the financials of the company. A company that has a *materiality threshold* (i.e., the minimum level of loss that matters significantly to the organization), for example, of $5 million that defines a material weakness in Sarbanes-Oxley compliance will prioritize its risks differently than a firm for which the materiality threshold is on the order of $500 million.[13]

With the risk analysis to use as a rationale, the road map to tailoring a book of policy derived from the international standards becomes a series of steps to take for each risk, to define the mitigating control, and to map that control to a control statement in, for example, ISO/IEC 17799:2005. Repeating this for all the risk categories and all the risk areas, one will have a foundation to demonstrate to any trading partner or outside party why the standard of care is relevant and appropriate to the company.

The next step is to help provide the translation of the high-level standard of care policy statements into action. Creating a set of implementation standards for the various controls described in the standard of care gives clear direction to all parties charged with ensuring that the needed controls are in place. The implementation standards also provide the foundation for a set of measurements and tests to determine that the controls are working as intended. This is the critical element that enables the executive

EXHIBIT 65.1 International Standards for Information-Security Governance

Basic set of internationally accepted standards includes	Useful in this context:
BS 7799	Replaced —see ISO/IEC 27001
ISO/IEC 17799:2005—Information technology. Security techniques. Code of practice for information-security management.	This universal standard provides a complete set of guidelines for an effective Information Security Management System (ISMS). It is essential guidance to help manage an effective information-security policy. It offers a common language and a common understanding to enable an organization to develop, implement, and measure effective security management practice, providing confidence in intercompany trading. ISO/IEC 17799 details a number of individual security controls, which may be selected and applied as part of the ISMS. ISO/IEC 17799, again based on a British standard, is scheduled to become ISO/IEC 27002 in a couple of years.
BS 7799-2:2005 (ISO/IEC 27001:2005) Information technology. Security techniques. Information-Security Management Systems. Requirements.	ISO/IEC 27001 specifies the requirements for the security management system itself. It is this standard, as opposed to ISO/IEC 17799, against which certification is offered. ISO/IEC 27001 has also been harmonized to be compatible with other management systems standards, such as ISO/IEC 9001 and ISO/IEC 14001. Organizations already certified under BS 7799- 2:2002 need to prepare for transition to ISO/IEC 27001 in order to meet its requirements.
BS ISO/IEC 13335-1:2004 Information technology. Guidelines for the management of IT security. Concepts and models for information and communications technology security management	This standard is useful to: • Define and describe the concepts associated with the management of IT security. • Identify the relationships between the management of IT security and management of IT in general. • Present several models that can be used to explain IT security. • Provide general guidance on the management of IT security.
IT Governance Institute's *Control Objectives for IT* (CoBIT®)	CoBIT® is a way to implement governance. It provides a tailorable set of controls.
ITIL—IT Infrastructure Library	ITIL® provides a cohesive set of best practice, drawn from the public and private sectors internationally. It is supported by a comprehensive qualifications scheme, accredited training organizations, and implementation and assessment tools. The best practice processes promoted in ITIL® support and are supported by the British Standards Institution's standard for IT service Management (BS15000). (OGC - ITIL)

team and the CISO to demonstrate that the controls are not only the right ones, but that they are working properly.

The CISO does not get involved in monitoring firewall rules or router configurations. This is work that is the domain of an operations security manager who is carrying out the standards established by the CISO. The CISO should review the reporting (once it is established) to ensure that the controls described for all Internet connectivity are in place and working, and to ensure that the handling of technical details is delegated to specialized technical staff and management.

65.3.2 Governance and Accountability. A recent research project completed by Booz Allen Hamilton reveals several factors that are changing the face of protecting enterprise assets, and ultimately adjusting the roles of the security profession. The information protection strategy is undergoing *convergence*, defined by ASIS International as "the identification of security risks and interdependencies between business functions and processes within the enterprise and the development of managed business process solutions to address those risks and interdependencies."[14] The role of the successful CISO is not only converging into an interdependent set of security specialties; it is also blending into business functions, intertwined into the fabric of decision-making processes throughout the life cycle of business strategy, plans, and execution.

Given the need for managing all the complexity recently introduced, security professionals must also adopt a different outlook—a more business-oriented than a protectionist position. Tim Mather, CISO of Symantec Corporation, describes it this way:

> For many information-security professionals, the urge is to promote information security—zealously. Many times too zealously. We often come to believe that security awareness equates to zealous promotion of information-security objectives, especially the deployment of (information) security technology—often at the "expense" of people, and policies and processes. However, that zealous promotion of information-security objectives tends to cloud our judgment as to the business considerations of the risks involved. Our information-security–colored glasses are polarized to security and tend to filter out business unit considerations. This leads to a loss of credibility with business unit personnel, hindering our ability to accomplish our information-security goals.
>
> Our challenge is to articulate our information-security objectives in terms of business risk that business unit personnel can understand and appreciate. That being said, it does not mean being "soft" on our objectives. We are not paid to install the "speed bumps" in enterprise hallways, but that articulation does mean translating information-security objectives into "business-speak." Only the combination of articulate translation and polite but firm emphasis on our information-security objectives will gain us the credibility that we need to accomplish our enterprise information-security goals. And that effort has to begin, and continue, at the top of the enterprise—with its executives.

Part of the reason security managers have had to take such hard positions in the past is that information-security practices have had a strong tendency to be in organizational silos: executive protection, IT, physical security, local business practices, policy management, HR investigations, fraud, and so on. Each of these practice areas is a specialty, to be sure, and will continue to require specialization. Independence can no longer be effective. Privacy, information security, data management, mobility, supply chain management, workforce management, human resources, facility management—the protection elements of these business domains, and so much more, must be coordinated to avoid weak links in the protection strategy that could prove disappointing if not devastating to a business, or to individuals.

With the need for an interdisciplinary approach at an executive level, with a broad reach across an enterprise, one might argue that the centralization of primary information security functions under a CISO would be one way to ensure that all information-security practice would be consistent and easily measured; on the contrary, the temptation to ensure the fidelity of an organizational mission through assembling affected functions under a single manager should be resisted. In a successful information-security strategy, it is far too easy for all enterprise management to think of security as *security's business* or as *the CISO's job*. In fact, the function of security should be emphasized as a business function, with accountability distributed in the business across all business unit executives, led by an executive-level security team. This has five advantages.

1. There will never be a security organization large enough to do the entire job of securing information assets in the new world. Security truly has to be everyone's job, in ways that are measured and tracked. Putting security accountability into the business units will leverage resources effectively.
2. Accountability is a good way to get someone's attention. Management is attentive to scorecards that are read at the top, with expectations for improvements that have been assigned to a responsible individual or group.
3. The interdisciplinary approach needed for the rapidly emerging challenges and competing priorities can best be addressed by an interdisciplinary team, including functions such as supply chain management, that have no reason to be merged into a security organization.
4. Being closely tied to the business allows for upstream integration into business processes, allowing security to be *built in* [versus an afterthought] when new business initiatives, strategies, and ventures are being designed.
5. The funding for proper protection of information assets is no longer a security problem or an IT problem, but it is a business problem, with the proper business visibility of what is and is not getting done, and what is the residual risk from funding decisions.

A simple policy-driven governance structure has been effective in many large-scale organizations to establish the necessary linkages with the business, to ensure consistency in strategy and approach, to gain executive buy-in on strategy through collective priority setting, and to gain visibility of progress and of unresolved issues on an ongoing basis. A basic governance structure can be tailored to fit corporate culture such as the number of representatives, and the layers of working groups. What is essential is the creation of a governance body to ensure corporate due diligence and to avoid conflicts of interest, such as are described further in Section 65.5.5.

A policy-driven approach to security and disaster recovery governance includes these components:

- Senior Leadership Team (C-level executives)
  - Authority for policy
  - Governance body
  - Program oversight
- Policy
  - Establishes accountability

  - Establishes programs and program authority
  - Establishes governance processes
  - Publishes under authority of the senior leadership team
- Principles
  - Five to six high-level statements at a descriptive level
  - Separate from policy
  - Establishes guidance for business unit standards
- Business Unit Team
  - Provides staff support
  - Facilitates governance process
  - Provides technical leadership for security across business units
  - Consists of representatives from each business unit plus audit
  - CISO is chair
  - Coordinates policy principles for senior leadership team approvals
  - Establish standards at the prescriptive level
  - Implements standards
  - Monitors effectiveness of implementation (provides metrics and reports to senior leadership team)
  - Coordinates key initiatives for security and recovery improvement across business units

In this structure, the CISO can function as an agent of change for driving broad, funded, prioritized initiatives with true business impact.

Joel Scambray, coauthor of *Hacking Exposed* and senior security strategist at Microsoft, sees the need for CISOs as change agents:

> Information security is now such a very broad topic, the role of the senior security professional has to move away from implementation of the security technology to the role of change agent. Network issues are diminishing: so much can go through port 80 inbound and then by proxy to the application, the themes in the attack community are shifting and will always continue to do so as technology evolves. Attackers will strike a business through the path of least resistance and that is going to be throughout the business, not just in the network and applications. It is essential to have business group accountability to address risks in all areas, for each executive to think about risk and ensure they are taking informed steps to bring it to within acceptable levels.
>
> The CISO has to move up in the corporate structure—it's a revenue-protecting job and has to be supported at the highest levels of the company to get the company to focus on the right risks.[15]

65.3.3 Roles and Responsibilities. In general terms, there are a set of roles and responsibilities that support accountability that have been shown effective in practice. Some would require substantial change in the organizational responsibilities within a company or agency. That ISO is the change agent to make things better; once the foundation for security governance is in place, this should be a discussion topic. Outside professional opinions can be solicited. Refer to the practices documented at Institute of Internal Auditors (IIA), Information Systems Audit and Control Association (ISACA), and the IT Governance Institute (ITGI). Evaluate them in the context of your company; you may not be able to implement all 10 of these principles for accountability

and reporting, but even some of them would demonstrate progress. The 10 principles for effective information-security control follow. For each item involving security, the CISO is the leader, or receives and acts on reports, or merely a participant and observer.

1. The CISO does not *own* IT assets but manages them on behalf of the business. Governance processes are in place for effective business management of IT.
2. An independent third party regularly reviews the implementation of each of these principles to verify control design and control effectiveness.
3. The expenditures on IT security are justified in terms of business value according to parameters established between the CISO and the business units.
4. IT security is actively monitored by an IT governance board consisting of IT, business management, and the CFO, and adjusted according to business needs.
5. Security, information assurance, and cybersecurity rules are tied to business rules in ways that are traceable, understandable, and agreed to by the business.
6. All IT security change is authorized by specifically designated IT management change boards.
7. Application development processes verify that applications perform only as intended, throughout the life cycle of the application, under the supervision of IT governance boards.
8. All IT security operations and processes are standardized, documented, and reviewed regularly for consistency by IT management and independent third parties.
   a. New processes are developed to accommodate business change.
   b. Existing processes are reviewed regularly for update, to accommodate business change. Consider having legal counsel review the documentation to determine if the records present unexpected legal risks or if they can provide legal advantage in any dispute over the standard of care.
9. All information systems assets (data, infrastructure, applications, processes, and services) have clear business owners with accountability to ensure:
   a. Assets are used only as intended.
   b. Assets are accessed only by those who are authorized according to defined business rules. (Access is defined as ability + opportunity.)
   c. Assets are available for use according to defined business rules.
10. Business and IT employees, contractors, vendors, and third parties have necessary documentation and training on a regular basis to carry out these principles in an effective, verifiable manner for the businesses in which they are engaged, and they have proof of training on a timely basis.

Generally, the business does not want to own IT. They see IT as a utility, the domain of IT staff, and too much work to understand. This is the major hurdle to be overcome in the change agent role of CISO. Remember, IT is the business. That is to say, IT today governs the systems that process the information that is the lifeblood of the business. Without the information, the business will stop. The logic applied to the information on behalf of the business should be owned by the business. It is worth the effort; once the business realizes its own accountability for the confidentiality, availability, integrity, ownership, and possession of business information, as well as

for information-security practices and improvements, it will enjoy a whole new level of interest in, and responsibility for, the enterprise.

65.3.4 Reporting. With the governance structure established, the ability exists to assign accountability to the right places in the organization. The senior leadership team should identify how management will be held accountable. This is not a choice of *if* but of *how*. Without clear accountability, governance and policy will mean nothing because they will not be implemented or enforced effectively. The primary tools for enabling executive accountability are reporting, monitoring, and metrics.

65.3.5 Monitoring. Monitoring in the context of the standard of care is the set of processes, human resources, and automated and manual tools needed to ascertain how well the controls established by the standard of care are functioning. Key to monitoring success is to monitor meaningfully—monitor to determine that the set of controls established is, or is not, working as intended. This is a function that can easily be coordinated and shared with internal audit, and care must be taken to ensure that the monitoring is appropriate for the applicable standard of care rules. Monitoring should be done with an objective degree of proof—hearsay is not adequate monitoring. Monitoring will usually need to be automated to handle the scale and frequency required in most information systems environments.

65.3.6 Metrics. Metrics have been a difficult challenge for all CISOs. The generic advice is this: Choose metrics for reporting that are essential to (a) give the business a set of key performance indicators and (b) make a difference in a needed control area. We often measure something in security just because we can. This is a mistake. Measuring the number of viruses or number of penetration attempts is not meaningful because they are high-volume certainties. The measurement system has to ensure that good results could not be achieved from failing to look; a decline in security incidents from outside attacks may be truly declining, based on 100 percent visibility of the problem, or the problem may have shifted to a space that is not monitored. In most cases it is probably, the latter, given the changing vectors in the network attack space.

65.3.7 Executive Visibility. For executive visibility, experience indicates that a CISO executive scorecard, published with the support of the CEO, CFO, COO, or other executive sponsor, will drive behavior according to the metrics that are chosen and reported. Reporting should be at least quarterly—more frequent would be desirable, but may be difficult to achieve. Quarterly frequency allows for continuity of program management in making security improvements that require executive support for implementation.

65.4 SUMMARY OF ACTIONS. In summary, the strategy for information security has to balance vulnerability management with a standard of care that is appropriate to all stakeholders. These stakeholders include business owners, trading partners, consumers, regulators, auditors, shareholders, and information owners whose membership in the interconnected community puts them at risk if any member has a breach. This requires coordination among diverse specialists and organizations. Security as a state of being is not feasible; due diligence to a standard of care is the approach CISOs should adopt.

The CISO has to develop a standard of care that answers to many demands. The Federal Information Security Management Act of 2002, EU Data Protection Act, Gramm-Leach-Bliley, HIPAA, Sarbanes-Oxley Act of 2002, Senate Bill 1386, BS 7799 certification, FIPS 199 and 200, as well as various NIST publications, and a host of local codes and statutes around the world greatly complicate the compliance function. It is best to adopt a standard-of-care approach that provides, based on international standards, a single comprehensive response to all queries. Expectations of records management, compliance with e-discovery, business continuance, disaster preparedness, IP and trade secret protections, sanctions for unfair information practices or advertising, and a burgeoning library of case law in matters related to security and privacy have become the drivers to which information-security professionals must respond. Clearly, information security has become a risk management role worthy of a C-level executive, working in a structured governance role to involve the business. Today, the transition from security director or even from today's CISO is not complete to the C-level role that the CISO name implies. The points made thus far indicate a real need to pursue this as a professional group and to adopt some generally accepted principles and standards so that businesses and agencies can enjoy the level of information protection that is required in this new world.

65.5 RECOMMENDATIONS FOR SUCCESS FOR CISOs. For all the interviews that I did at RSA in 2006, only one quote made it into the press. Under the title "Microsoft CISO Has a Sense of Humor," the reporter indicated that I had answered all their detailed questions with serious thoughtfulness. But when they asked me what would be the best advice to someone who aspires to be a CISO, they quoted, "Have a stiff drink until the feeling goes away." I actually said, "Have a lie-down somewhere," but no matter—the quip hit the papers.

It was seen as funny, but there is also a real need to consider the demands made on an individual who aspires to a C-level role in security, where there is so much at stake. Many technical security managers are chosen to fill the role but are ill equipped to transform into executive management. Even worse, a CISO who does understand the executive role and moves to perform it in an organization where CISO roles are not understood can easily fail. To quote Machiavelli:

> Let it be noted that there is no more delicate matter to take in hand, nor more dangerous to conduct, nor more doubtful in its success, than to set up as a leader in the introduction of changes. For he who innovates will have for his enemies all those who are well off under the existing order, and only lukewarm supporters in those who might be better off under the new.[16]

In this section, several pointers based on input from various CISOs are offered to help you assess whether an organization is ready for a CISO and whether you are ready. The areas include:

- Education and experience
- The *culture* of security in the business
- The alliance with corporate and outside counsel
- The partnership with internal audit
- The tension with IT and dealing with the potential conflict of interest
- Organizational structure
- Responsibilities and opportunities outside of CISO internal responsibilities

65.5.1 Education and Experience. First and foremost, do you have appropriate background to be able to lead with wise counsel and appropriate judgment? Do you feel you are better at tactical detail, or do you prefer the big picture? In a panel during a plenary session at RSA 2005, LJ (Lisa) Johnson, CISO of Nike, related that her MBA degree was very helpful in her role as CISO, and this sentiment echoes across others in similar positions; however, successful leaders in this space do not all have advanced degrees—life experience is a significant contributor. Reading the wealth of books on various aspects of business management, information systems, and information security is also a good way to expand one's background. Conferences are beginning to focus on the role of CISOs as the role is described in this chapter, but, for the most part, conferences are best at keeping up with technical or auditing trends.

One way to evaluate the background that is necessary, whether it be by experience, reading, or the classroom, is to evaluate the various stakeholders with whom a CISO must effectively communicate to one degree or another. Internal audit, legal counsel, business executives, and IT staff are the core stakeholders. Within the business units, finance, supply chain, and HR are areas with which the CISO will have frequent discussions. Do you know your company's value chain? Do you understand the major business processes? Have you read the company's disclosure statements, and annual report? Do you understand the major cost concerns and the revenue streams? What are the key risks, outside of security, that occupy the executive team's time? If you know the answers to these questions, you are in the minority of security professionals. If you do not create your own action plan to network with key individuals, then read, take a course, and create a career plan that will give you experience in these areas. Do not be afraid to admit that you are learning and interested in getting more information from the experts in a particular area.

Whether your education is a BS in computer science, an MBA, or high school, you can get a great education on the job in preparation to taking a leadership role in information security. Evaluate what it will take for success, and then make your plan to get that information on the job or in school.

Each company is unique, so each CISO job will have unique elements built on the basic CISO job description. A word of warning: As a new CISO, take the time to absorb the culture before launching major changes or initiatives. Remember Machiavelli—people need to know you care before they care what you know.

65.5.2 Culture of Security in the Business. The culture of security can be very difficult to ascertain. In a centralized security model that lends itself to top-down management, one at least knows where to look. In a company where security may be part of the company's core business, and many think of themselves as CISOs, the role is much more challenging. The culture can be defined as the set of attitudes toward accepting direction, allowing time and resources to be used in order to put proper controls in place. The next points will help you take a pulse check on culture and could be used for some due diligence prior to accepting a CISO role:

- **Risk appetite.** What is the materiality threshold (minimum significant loss) for risk management? It is important to understand this so that you can properly ascertain the importance of the issues you will encounter and thereby know whether they require an FYI or escalation as an urgent matter.
- **Cultural norms and attitudes.** Is the workforce dynamic, high rate of job turnover, low threshold of tolerance for change, autonomous? You would work

with this kind of workforce and the culture it engenders much differently than a company with a by the rule-book, policy-driven bureaucracy. There is no right answer, but you need to know how people receive security direction. Most people favor a carrot over a stick; it is a question of what kind of carrot.

- **Relevant regulations and statutes.** You will need an inventory and at least a basic understanding of the way regulations and statutes affect information security. If the company is multinational, this requires an understanding of local jurisdictions, local statutes, and local regulations as well. Find the people who can share with you what they know about this area.
- **Influence and awareness of the court of public opinion.** The court of public opinion is essentially a form of reputation risk and the impact that it can have on the organization. A decision about security, such as fixing a security bug in an online application, may not be governed by specific laws, but it could definitely create the perception of poor security practices in a way that influences other security issues that the company may have in the public eye. This area may be a major influence for highly visible companies, or less of a concern for others.
- **Influence of risk to reputation.** With consumer confidence at an all-time low for Internet security, this is an element of the company's risk profile that it is essential to define explicitly as much as possible. Would the company go so far as to achieve full transparency of its security controls and privacy statements as an approach to building trust as a competitive advantage? Or is it willing to take more reputation risk by obscuring the privacy statement because the security controls are not quite what they should be? As you discover security issues, knowing the company's posture relative to transparency and associated reputation risk will be useful to determine what to escalate.
- **Who else manages various aspects of risk?** Make it a point to build and maintain an active network of peers in these areas:
  - Financial
  - Legal
  - Corporate strategy and planning
  - Marketing
  - Other areas of security (consulting, product, customer support)
- **How is your role seen in relationship to those managers?** Who else should you be talking to? Keep asking the questions.

65.5.3 Alliance with Corporate and Outside Counsel. There are a few roles that merit special mention. Corporate counsel and outside counsel are becoming increasingly interdependent with the information security team. Historically, company policy, investigations, contracts, and incidents have all (or should be) coordinated with counsel. New developments in records management, e-discovery, third-party management, privacy, and case law relating to information-security breaches and losses have created a new role for counsel in defining protection of consumer and employee data and company IP. Counsel will be necessary to navigate the rule of law as it pertains to information, information systems, and associated business practices. This is an organization with which the CISO needs to be close.

65.5.4 Partnership with Internal Audit. While it is not uncommon to find IT personnel who assume that what audit does not know cannot hurt them, this attitude

is a grave mistake. Internal audit is another special relationship for the CISO. A close IT–audit partnership is essential for these key reasons, and it should be the responsibility of the CISO to ensure that this partnership is in place and working smoothly:

- The rules for standard of care require demonstration that the controls are working effectively as designed. There is no better organization to provide clarity on what that entails than internal audit.
- Internal audit can collaborate with information security on the standard of care monitoring and reporting, thus effectively extending information protection resources.
- The CISO is accountable to ensure that control processes and policies that pertain are in place and working well. Internal audit should have full access to information-security processes to ensure compliance and to provide independent assurance against potential conflict of interest.
- By cultivating a close partnership (yet maintaining the arm's-length relationship), the CISO and an organization can exchange information, set priorities for the audit and security improvement programs, and provide support for each other on core issues to escalate.

65.5.5 Tension with IT. Many CISOs report to IT, and therein lies a potential for conflict of interest, or the appearance of conflict of interest. This is an area of increasing concern. The issue resides with the CISO needing to identify the security concerns and strategies for the company, including scope that extends beyond IT control, but doing so may be perceived as reflecting poorly on the IT management chain. Also, having a CISO report to the CIO creates an impression that the CISO is the IT security manager. Further, IT budgets are often calculated on the basis of run rates, assuming that it is an operational utility. Funding all the security effort out of an IT budget is a mismatch in two ways: Security is not a run-rate type of function, and taking what amounts to 10 percent or more of the IT budget for security is generally a significant hardship on other parts of IT that are critical to the business.

To address the conflict of interest concern, one could move the CISO role out of IT and have it reporting at the CFO, COO, or CEO level. The CISO in this regard would be a peer of the CIO. This has an advantage of removing the conflict of interest and the impression of IT security manager, but it also introduces the possibility of the CISO losing touch with the IT organization. Another alternative, if the CISO reports to the CIO, is to establish the governance structure and provide dotted-line relationships to the senior leadership on the governance board. Internal audit can monitor the relationships and processes to ensure that conflicts of interest are not developing. Politics being what they generally are, the latter choice seems the most difficult to implement well.

65.5.6 Organizational Structure. Organizational structure for the CISO will, of course, be heavily influenced by the company. We have already discussed some independent factors that argue against the CISO reporting to the CIO. Companies are moving away from the traditional model of having the senior information-security manager, or in some cases CISO, report directly to the CIO, opting instead to have the CISO report to the CEO, COO, or, in some cases, the CFO. In addition, a cross-functional, dotted-line relationship to key stakeholders may be required for adequate reporting and oversight.[17]

Aside from the question of to whom should the CISO report, what should the CISO manage? How should that be structured? The CISO's role as an executive should be

focused on governance, policy management, and compliance monitoring and reporting. The CISO should establish parameters for IT security operations, information-security investigations, forensics and incident handing, identity and access management, business continuance, records management, and e-discovery but need not necessarily handle the day-to-day oversight of those functions. The CISO could also be responsible for physical security and executive protection or have oversight of those roles with other security managers.[18] The role has grown to become more than any single person can track if all these functions report directly to the CISO. It is incumbent on the CISO to set the strategy and the structure for all these related security functions and to monitor their progress against predefined performance objectives.

65.5.7 Responsibilities and Opportunities outside of CISO Internal Responsibilities. In this chapter, we have described CISOs as change agents and strategists, adopting a standard of care and strategies that both implement that standard of care and demonstrate due diligence to it. Such individuals have much to contribute to the profession, to the community at large, to the technical community, and even to the definitions being established by governing bodies at the local, state, federal, and international levels.

It would be unreasonable to try to prescribe what CISOs should do outside their direct internal responsibilities. Suffice it to say that CISOs should subscribe to a professional ethic to share what they have experienced, to codify security practice, and to bring about a better understanding of the problem space, so that defined problems can be solved with best practices that ultimately become standard protocols. Make it a practice to participate in professional organizations, to write, and to speak to community, professional, and trade organizations. Get the word out. Work to eliminate confusion. Define the role. There is much at stake for those who choose to wear this mantle—in the words of Theodore Roosevelt, "It behooves every man to remember that the work of the critic is of altogether secondary importance, and that in the end, progress is accomplished by the man who does things."

65.6 CONCLUDING REMARKS. CISOs are strategic, executive agents of change for the protection of information that is the lifeblood of critical infrastructures and private enterprise. The realm of disciplines that CISOs must manage is expanding as the scope of security has broadened—an interdisciplinary approach that involves all aspects of the business to provide digital, physical, and personnel security is required with input from vital stakeholders.[19]

This role is different from that of technical security management in the near past, and involves tools that are different. The complexity and sheer insecurity of the interconnected ecosystem requires that we adopt a standard of care and that we be able to demonstrate due diligence to that standard of care—this is the mission of CISOs.

CISOs, to be successful, depend on strong governance, roles, and responsibilities; management accountability; and strong reporting practices (including monitoring, metrics, and executive visibility) to be successful. CISOs would necessarily provide oversight, direct or indirect, of functional roles such as technical security management functions.

CISOs are a rare breed and need to make their contributions to their respective organizations as well as to the larger set of professional, community, and industry groups that need guidance, clarity, and judgment in their movement forward to a trusted, interconnected ecosystem.

The challenge is huge—that is why CISOs like it. It is ever-changing, touches every part of business and technology, and is ultimately a people job. It is a very difficult job and, under the right circumstances, is a very rewarding one. It is a high position of trust and responsibility, and it finally has come into its own.

65.7 FURTHER READING

Brotby, W. K. *Information Security Governance: A Practical Development and Implementation Approach*. Wiley, 2009.

Brotby, W. K., and G. Hinson. *PRAGMATIC Security Metrics: Applying Metametrics to Information Security*. Auerbach Publications, 2013.

Fitzgerald, T. *Information Security Governance Simplified: From the Boardroom to the Keyboard*. CRC Press, 2011.

Fitzgerald, T., and M. Kraus. *CISO Leadership: Essential Principles for Success*. Auerbach Publications, 2007.

Harkins, M. *Managing Risk and Information Security: Protect to Enable*. Apress, 2012.

Kouns, J., and B. L. Kouns. *The Chief Information Security Officer*. IT Governance Publishing, 2011.

Oberlaender, M. S. *C(I)SO—And Now What?: How to Successfully Build Security by Design*. CreateSpace Independent Publishing Platform, 2013.

65.8 NOTES

1. Thomas L. Friedman, *The World Is Flat: A Brief History of the Twenty-First Century* (Farrar, Straus, and Giroux, 2005).
2. Eugene Spafford, "Information Security: Insanity Rules," AusCERT2006, AusCERT, Royal Pines Resort, Gold Coast, Australia, May 24, 2006.
3. Larry Greenemeier, "VA Had Many Security Warnings before Its 26.5 Million-Person Breach," *Information Week*, www.informationweek.com/va-had-many-security-warnings-before-its/188500807
4. T. Cooper, "NZ 'definitely not for sale': Ocker* tries to flog God's Own on eBay" *The Register*, May 12, 2006, www.theregister.co.uk/2006/05/12/nz_auctioned_on_ebay
5. Bank for International Settlements, "Basel II: Revised International Capital Framework" (2006), www.bis.org/publ/bcbsca.htm
6. Scott Berinato, with Lorraine Cosgrove Ware, "Six Secrets of Highly Secure Organizations," *CIO Magazine*, September 15, 2005, www.coresecurity.com/content/six-secrets-of-highly-secure-organizations
7. P. Condit, Keynote Speech, International Information Integrity Institute Forum 25, The Boeing Company, Seattle, WA (January 1995).
8. United States, Office of the National Counter-Intelligence Executive, *Annual Report to Congress on Foreign Economic Collection and Industrial Espionage—2004* (Washington, DC: ONCIX, 2005). www.fas.org/irp/ops/ci/docs/2004.pdf
9. Donn Parker, "Making the Case for Replacing Risk-Based Security," ISSA Journal (May 2001): 6–10.
10. P. Schwartz, *The Art of the Long View: Planning for the Future in an Uncertain World* (Doubleday, 1991).

11. P. W. Parfomak, "Guarding America: Security Guards and U.S. Critical Infrastructure Protection," Congressional Research Service Report for Congress, November 12, 2004, Order Code RL32670; http://ftp.fas.org/sgp/crs/RL32670.pdf
12. Ian Davis and Elizabeth Stephenson, "Ten Trends to Watch in 2006," *McKinsey Quarterly*, http://blog.jackvinson.com/archives/2006/01/19/the_mckinsey_quarterly_ten_trends_to_watch_in_2006.html
13. The concept of materiality threshold is widely used in accounting. For example, see the "Accounting Terminology Guide" from the New York State Society of CPAs, www.nysscpa.org/prof_library/guide.htm
14. Booz Allen Hamilton, "Convergence of Enterprise Security Organizations," ASIS (Alliance for Enterprise Security Risk Management) International, November 8, 2005, http://iris.nyit.edu/~kkhoo/Fall2008/755/ConvergenceEnterpriseSecurityOrg_08Nov05.pdf
15. Stuart McClure, Joel Scambray, and George Kurtz, *Hacking Exposed*, 5th ed. (Emeryville, CA: McGraw-Hill Osborne, 2005).
16. Niccolo Machiavelli, The Prince (1513), Chapter VI, Para 5, www.gutenberg.org/files/1232/1232-h/1232-h.htm
17. Scott Berinato, with Lorraine Cosgrove Ware, "Six Secrets of Highly Secure Organizations."
18. See, for example, Fran Howarth, "The Convergence of Physical and IT Security," IT-Director.com, September 11, 2006, www.it-director.com/business/regulation/content.php?cid=8743
19. Thomas L. Friedman, *The World Is Flat*.

CHAPTER 66

DEVELOPING SECURITY POLICIES

M. E. Kabay and Sean Kelley

66.1 INTRODUCTION 66·2

66.2 COLLABORATING IN BUILDING SECURITY POLICIES 66·2

66.3 PHASE 1: PRELIMINARY EVALUATION 66·3

66.3.1 Introduction to the Study 66·4
66.3.2 State of Current Policy 66·5
66.3.3 Data Classification 66·5
66.3.4 Sensitive Systems 66·5
66.3.5 Critical Systems 66·6
66.3.6 Authenticity 66·6
66.3.7 Exposure 66·6
66.3.8 Human Resources, Management, and Employee Security Awareness 66·6
66.3.9 Physical Security 66·7
66.3.10 Software Development Security 66·7
66.3.11 Computer Operations Security 66·8
66.3.12 Data Access Controls 66·9
66.3.13 Network and Communications Security 66·9
66.3.14 Antimalware Measures 66·10
66.3.15 Backups, Archives, and Data Destruction 66·10
66.3.16 Incident Response 66·11
66.3.17 Business Resumption Planning and Disaster Recovery 66·11

66.4 PHASE 2: MANAGEMENT SENSITIZATION 66·11

66.5 PHASE 3: NEEDS ANALYSIS 66·12

66.6 PHASE 4: POLICIES AND PROCEDURES 66·12

66.7 PHASE 5: IMPLEMENTATION 66·13

66.7.1 Upper Management 66·13
66.7.2 Technical Support 66·14
66.7.3 Lower-Level Staff 66·14
66.7.4 Other Technical Staff 66·14

66.8 PHASE 6: MAINTENANCE 66·14

66.9 CONCLUDING REMARKS 66·14

66.10 FURTHER READING 66·15

66.11 NOTES 66·15

66.1 INTRODUCTION. This chapter reviews methods for developing security policies in specific organizations. Some of the other chapters of this *Handbook* that bear on policy content, development, and implementation are listed next:

- Chapter 23 provides an extensive overview of physical security policies.
- Chapter 25 discusses local area network security issues and policies.
- Chapter 39 reviews software development policies and quality assurance policies.
- Chapter 44 presents resources and standards for creating effective security policies.
- Chapter 49 looks at methods for enhancing security awareness.
- Chapter 45 provides guidance on employment policies from a security standpoint.
- Chapter 47 makes explicit recommendations about operations management policies.
- Chapter 48 reviews specific recommendations for email and Internet usage.
- Chapter 50 presents concepts and techniques from social psychology to make security policy implementation more effective.
- Chapter 52 discusses the policies that apply to application design.
- Chapter 72 looks at censorship and content filtering on the Internet.

66.2 COLLABORATING IN BUILDING SECURITY POLICIES. Policies are the foundation of effective information security, but the task of policy creation is complicated by human and organizational resistance. Technology alone does not work. In changing human behavior, rationality and substance are not enough: The *process* of development affects how people feel about policies and whether they see these rules as needless imposition of power or as an expression of their own values.

Security is always described as being everyone's business; however, in practice, security *interferes* with everyone's business. For example, network managers work hard to make networks user friendly. They do everything they can to make life easier for users; they provide network access routines with a graphical user interface, client/server systems with hot links between local spreadsheets and corporate databases, and a gateway to the Internet for their users. Superficially, one might think that implementing network security would simply involve defining access controls, applying encryption, and providing people with handheld password generators. Unfortunately, as discussed in Chapter 50, security policies offend deep-seated self-conceptions. People form close-knit work groups in which they trust each other; they do not lock their desks when they leave them for a few minutes, so why should they obey the network security policy that dictates locking their sessions? They even lend people car keys in an emergency; why should it be such a terrible breach of security to lend access codes and passwords to trusted colleagues in an emergency?

Security policies challenge users to change the way they think about their own responsibility for protecting corporate information. Attempting to impose security policies on unwilling people results in resistance, both because more stringent security procedures make people's jobs harder and because people do not like being told what to do—especially by security officials perceived as being outside the chain of command.

The only approach that works in the long run is to present security to everyone in the organization in a way that causes recognition that each one, personally and professionally, has a stake in information protection. Security managers, to be successful,

must involve employees from throughout the enterprise in developing security policies. Users must justifiably feel that they own their security procedures; employees with true involvement in the policy development process become partners, rather than opponents, of effective security.

66.3 PHASE 1: PRELIMINARY EVALUATION. Studies of the extent to which information-security policies are in place consistently show that relatively few of the respondents have adequate policies in place. For example, the 2007 Global Security Survey run by *CIO Magazine* and PricewaterhouseCoopers was based on interviews and questionnaires involving 7,200 executives, security professionals, and technology managers "across all industries and more than 100 countries" about their organization's security and privacy policies and practices. According to the summary on page 2 of the report:

1. Only 6 out of 10 respondents worldwide (57 percent) say their organization has an information-security strategy in place. Only another 13 percent, however, consider putting one in place a "top priority."
2. Almost half of all respondents worldwide (48 percent) say their organization does not actively engage both business and information technology decision makers in addressing information security.
3. More than 7 out of 10 security managers, administrators, and technicians worldwide believe that their organization's security policies and spending can be improved.[1]

In what follows, it is assumed that a specific officer or manager (or group of officers or managers) in the enterprise has taken on the task of developing security policies. The group will be called the *policy development group.*

Before attempting to formulate policies, the policy development group needs formal authorization to use corporate resources in such a project. It should not be too difficult to obtain a short memorandum from top management to everyone in the organization that lays out the reasons for asking for their time and energy in gathering information about the current state of security. Such authorization and continuing top-level support are essential tools in convincing people to cooperate with the policy development group.

In the absence of existing or adequate security policies, a preliminary inventory is the first step in providing upper management with the baseline information that will justify developing a corporate information-security policy. The preliminary evaluation should be quick and inexpensive—perhaps days of work by a few people. There is no point in wasting time in expensive detail work before getting approval, support, and budget from upper management.

The goal of the preliminary evaluation is to ask the people who work with information resources what they believe are their most important-security needs. Even though they may not be conscious of security as a distinct need, in practice, employees and managers do have valuable insights that transcend theory and generalizations. Data entry clerks may tell the security staff about security violations that no one else has observed or even thought about; for example, they may observe that a bug in a particular program makes the previous operator's data entry screen available for unauthorized entries when the shift changes and a new operator sits at the same terminal.

The policy development group should work closely with human resources (HR) personnel in developing the research instruments for interviewing staff. HR members

are likely to know the key managers to contact in each department. The managers have to be convinced to support the effort so researchers can interview willing staff. Some of the HR people are likely to have the professional skills and experience required to provide accurate and cost-effective evaluations of beliefs, attitudes, and behavior affecting security. They may be able to help construct unbiased questionnaires, organize focus groups, and guide interviews.

However, if the security staff and the HR staff are not confident about being able to handle this preliminary data collection, the policy development group should see if it can obtain authorization to hire a consultant with proven expertise in collecting and analyzing social attitudes. The policy development group might want to discuss such a study with a firm specializing in security audits and organizational analysis. If no one knows where to start looking for such resources, the policy development group can contact information security associations, security magazines, security Websites, and local universities and colleges to ask for suggestions.

These key issues should be part of the preliminary study:

- Introduction to the study
- State of current policy
- Data classification
- Sensitive systems
- Critical systems
- Authenticity
- Exposure
- Human resources, management, and employee security awareness
- Physical security
- Software development security
- Computer operations security
- Data access controls
- Network and communications security
- Antimalware measures
- Backups, archives, and data destruction
- Business resumption planning and disaster recovery

The next sections suggest some typical questions that would be helpful in gathering baseline data about the current state of security. All these questions (and more site-specific topics) should be asked of all the respondents in the preliminary evaluation. Applicable questions are not necessarily repeated in each section; instead, questions in the earlier parts of this list may be adapted for use in later sections. These suggestions are not intended to limit creativity but rather to stimulate development of more questions that would be particularly useful for a specific enterprise.

66.3.1 Introduction to the Study. Employees may perceive many of the questions as threatening. The preamble or introduction to the study, whether it is by survey or by interviews, should make it clear that this is *not an audit* and that its purpose is to

establish the framework for an appropriate set of security policies—policies suited to the needs of the organization and its stakeholders. The information should be anonymized so that no person will be targeted for reprisal if the study discovers problems. Every effort should be made to reassure employees that the study is designed to learn about the facts of security with a view to improvement, rather than a search for culprits who will be punished.

66.3.2 State of Current Policy. The questions that follow not only gather baseline information about security policies but also determine whether employees have any idea about who is responsible for formulation of those policies.

- Does the enterprise have any security policies at all?
- Who developed them: an individual? a group?
- Where and how are the security policies available (paper, electronic)?
- When were the policies last updated? Last disseminated?
- Who, if anyone, has explicit responsibility for maintaining security policies?
- Who implements security policy at the enterprise level?
- To whom does the chief information security officer report within the enterprise?
- Who monitors compliance with security policies, standards, and compliance?

66.3.3 Data Classification. Questions to ask include:

- Are there levels of security classification that apply to your work? If so, what are they called?
- Are there rules for determining whether information you handle should be classified at a particular level of confidentiality?
- Are documents or files labeled to show their security classification?
- What is your opinion about the value of such classification?
- Do people in your group pay attention to security classifications?
- Do you have any suggestions for improvement of how data are classified?[2]

66.3.4 Sensitive Systems. The questions in this section focus on information that ought to be controlled against unauthorized disclosure and dissemination.

- In your work, are there any kinds of information, documents, or systems that you feel should be protected against unauthorized disclosure? If so, name them.
- How do you personally protect sensitive information that you handle?
- How do others in your department deal with sensitive information? No names, please.
- To your knowledge, have there been any problems with release of sensitive information in your department?
- Do you have any suggestions for improving the handling of sensitive data in your area?

66.3.5 Critical Systems. The questions in this section focus on information that requires special attention to availability and correctness.

- In your work, are there any kinds of information, documents, or systems that you feel are so critical that they *must* be protected against unauthorized modification or destruction? If so, name them.
- Are there any special precautions you use, or know of, to safeguard critical data in your area?

66.3.6 Authenticity. Questions to ask include:

- Do you know of any cases in which anyone has used someone else's identity in sending out messages such as letters, faxes, or email? If so, were there any consequences?
- Does everyone in your group use digital signatures on electronic documents?
- Does anyone in your group make or use unauthorized copies of proprietary software? If so, do you think there is any problem with that?

66.3.7 Exposure. Questions to ask include:

- What are the worst consequences you can realistically imagine that might result from publication in the newspapers of the most sensitive information you control?
- What might happen, in your opinion, if key competitors obtained specific confidential information that you use or control in your area?
- Can you estimate monetary costs associated with the scenarios you have just described?
- What would be the worst consequences you can foresee if critical information you work with were to be altered without authorization, or through accidental modification?
- What might happen if you could not access critical information quickly enough for your work?
- Can you estimate the costs of such breaches of data integrity and data availability?
- Could there be trouble if someone forged documents in your name, or in the enterprise's name? Can you sketch out some scenarios and associated costs resulting from such breaches of authenticity?

66.3.8 Human Resources, Management, and Employee Security Awareness

- As far as you know, who is responsible for developing security policies?
- Do you know where to find the security policies that apply to your work?
- When, if ever, did you last sign any documents dealing with your agreement to security policies?
- Who is responsible for monitoring compliance with security policy in your workgroup? In the enterprise as a whole?

- Have you ever received any training in security policies? If so, when was the last time?
- Have you ever seen any written materials circulating in your workgroup that discuss information security?
- Do you think of protecting corporate information as one of your official responsibilities?[3]

66.3.9 Physical Security

- Does anyone check your identity when you enter the building where you work? Always?
- Are there any electronic access-control systems limiting access to your work area? What are they?
- Do people hold a secured door open to let each other into your work area? Do you let people in after you open a secured door?
- Have you ever seen a secured door into your area that has been blocked open (e.g., for deliveries)?
- Do people leave your work area unlocked when everyone leaves?
- Do staff members wear identity badges at work? Are they supposed to? Do you wear your badge at work?
- Do visitors wear badges?
- Have you ever seen strangers in your area who are not wearing visitor badges?
- What would you do if you saw a stranger in your area who was not wearing a visitor's badge?
- Do you lock any parts of your desk when you leave your workspace?
- What would you do if you heard the fire alarm ring?
- Where is the nearest fire extinguisher?
- Who is the fire marshal for your floor?
- What would you do if someone needed emergency medical attention?
- Is there an emergency medical station in your area or on your floor?
- Do you know who is qualified in cardiopulmonary resuscitation (CPR) in your group or on your floor? Do such people wear identifying pins?
- Have you had recent training in what to do in the event of an emergency? Have you been trained in how to evacuate the building?
- Is there anything that comes to mind that you would like to see to improve physical security and safety in your work area?[4]

66.3.10 Software Development Security. These questions would be asked only of the software development team:

- Are there any security policies that apply to your work? What are they?
- Have you ever discussed security policies in your group?
- Is security viewed positively, neutrally, or negatively in your group? And by yourself?

- Do you and your colleagues discuss security during the requirements analysis and specification phases when developing software?
- How do you see quality assurance as part of the development process?
- Do you use automated software testing tools?
- Do you use automated version control software? If not, how do you exercise control?
- How do you document your systems?
- Do you think that your source code is adequately protected against unauthorized disclosure and modification?
- What is your opinion about Easter eggs (unauthorized code for an amusing picture or game)?
- Could anyone plant an Easter egg or a logic bomb (unauthorized, harmful functions) in code being developed in your group?
- Have you ever seen an Easter egg or a logic bomb in code from your group? Did it get through to production?
- Can you think of ways you would like to see better security in your work?[5]

66.3.11 Computer Operations Security. These questions would be asked only of the computer operations team:

- How long do you wait after initial release before installing new operating system versions on your production machines?
- How do you put new software into production?
- Can development personnel access production software? Production data?
- How do you handle problem reports? Do you have an automated trouble-ticket system?
- Can people from outside the operations group enter the operations center?
- Are contractors, including repair technicians, allowed to circulate in operations without being accompanied?
- Does the cleaning staff ever circulate within the secured areas of operations without operations staff present?
- Are system components labeled?
- Is there an emergency cutoff switch for main power to the entire data center? Does it include air conditioning?
- Are there uninterruptible power supplies for critical components of your systems?
- Do you keep records of system downtime? What is your downtime over the last three months? The last year?
- What accounts for most of the downtime? Have remedies been implemented?
- Who monitors system resource utilization? Are there automated reports showing trends in disk space usage? CPU utilization? Network bandwidth usage?
- What improvements in security would you like to see in operations?[6]

66.3.12 Data Access Controls

- Do you have to identify yourself to the computers and networks you work with?
- Do you have a user name (ID) that no one else shares?
- Are you required to use a password, passphrase, or personal identification number (PIN) as part of your routine when starting to use your computer?
- Have you ever shared your unique user ID and password or PIN with someone else? Or have you borrowed someone else's user ID and password to get some work done? If so, how often does this happen?
- Do you use a token, such as a physical key or a smart card, to prove who you are to the computer system? If so, have you ever lent or borrowed such tokens? What for? How often?
- In your work, are there any limitations on the data you are allowed to see, modify, add to, or delete?
- Are there data you can see but not change?
- Do you use encryption on any of the data you work with?
- Do you or members of your group use laptop computers? If so, do you encrypt sensitive data on the disks of those portable systems?
- Do you or anyone in your group take work home? If so, do you put corporate data on your own, personal (noncompany) computers? Does anyone else have access to those computers? Are there any controls on accessing corporate data on the home computers?[7]

66.3.13 Network and Communications Security. Most of the next questions would be appropriate only for network managers, administrators, and technicians. However, some of the questions are suitable for everyone.

- As a user, do you know what the rules are about using your employer's email system and Internet access?
- Do you know anyone who regularly violates system usage restrictions? No names, please.
- Have you ever seen pornography on corporate systems? Child pornography? Racist and other objectionable materials? If so, did you know what to do? And what did you do?
- Has anyone ever discussed rules for secure email with you? Do you know how to encrypt sensitive messages? Do you ever encrypt messages?
- As a network manager, do you have up-to-date network diagrams, or can you produce them on demand?
- Do you know which services are running on your Internet-connected systems? Are all of the running services needed?
- How do you determine which patches are appropriate for installation on your systems? How often do you check? Who is responsible for managing patches? How long does it take between notification of a vulnerability and installation of an appropriate patch?

- Does your security architecture include firewalls? If so, what determines the security policies you instantiate in the filtering rules?
- Do you have egress filtering enabled on your firewalls?
- Do you have intrusion detection systems? If so, who responds to apprehended intrusions? How are the responsible people notified of an intrusion?
- What are the procedures for responding to an intrusion?
- If your organization uses passwords, how do you handle requests for new passwords?
- Do you have centralized remote-access controls?
- Do remote users use virtual private networks (VPNs) to access corporate systems from outside the firewalls?
- Are your users supposed to use encryption for sensitive email that traverses the Internet? Do they? How do you know?
- Do your users apply digital signatures to all communications?
- Are your Web servers protected against intrusion and vandalism?
- Have you kept sensitive information off your Web servers?
- Do you encrypt all sensitive information stored on your Web servers?
- How long would it take you to recover a valid version of the Website if it were destroyed or vandalized?
- Do your telephone voicemail boxes have unique, nonstandard passwords? How do you know?
- How do you find out if an employee is being fired or has resigned? How long does it take between termination of employment of such an employee and deactivation of all system and network access?

66.3.14 Antimalware Measures

- Do you and all of your users have antimalware products installed on every workstation?
- How often are antimalware products updated? How are they updated?
- How long does it take for all vulnerable systems to be brought up to date?
- Do you or your users open unexpected, unsolicited email attachments?[8]

66.3.15 Backups, Archives, and Data Destruction

- How often do you do backups of your electronic data?
- Where do you store backup media? Are current copies retained off site as well as on? How do you know which media to use to restore a specific file?
- How long do you keep different types of backups? Why?
- How do you prevent unauthorized access to backup media?
- If you keep data backups for several years, how do you ensure that the old media will be readable and that the data developed for old applications will be usable?
- How do you dispose of magnetic and optical storage media after their useful life is over? Are the discarded media readable?

- Do you make backup copies of paper documents? Where are these copies kept? How would you locate a specific document you needed?
- How long do you keep various types of papers? Why?
- When you dispose of paper documents, does their content influence how they are destroyed? How do you dispose of sensitive paper documents?[9]

66.3.16 Incident Response

- Do you know what to do if you see a security problem such as strangers without badges in secured areas, threats of violence arriving by email, damaged company Web pages, bad data in files you use, missing files or data, or forged messages?
- Is there a specific team dedicated to responding to computer security breaches?
- Are you part of a computer security incident response team (CSIRT)? If so:
  - Is this a long-term assignment?
  - Do you have documentation and training on how to respond to specific types of incidents?
  - How are out-of-hours incidents handled?
  - Does your CSIRT hold practice drills?
  - Are there incident postmortems?
  - Do you have relations with local law enforcement officials? If so, which levels (municipal, state, federal)?
  - Have you ever experienced pressures from management to suppress incident reports of criminal activity to law enforcement?[10]

66.3.17 Business Resumption Planning and Disaster Recovery

- Do you have business resumption planning (BRP) or disaster recovery plans (DRP)? If so, where are they kept?
- Who is responsible for keeping BRP and DRP up to date?
- Have you ever participated in a BRP or DRP testing? If so, how long ago was the last one? When is the next scheduled test?
- During BRP and DRP tests, does anyone use video cameras or tape recorders to keep track of critical steps in the recovery?
- After a test, have you participated in analyzing the results of the tests to improve the plans?[11]

66.4 PHASE 2: MANAGEMENT SENSITIZATION. Support from upper management is essential for further progress. The goal in this phase is to get approval for an organization-wide audit and for a policy formulation project. In conjunction with the rest of the information-security project team, the responsible managers should plan on a meeting that lasts no more than one or two hours. The meeting should start with a short statement from a senior executive about the crucial role of information in the organization's business.

Professional aids, such as management-oriented training videos, are helpful to sensitize managers to the consequences of poor information security. For an up-to-date list of such videos, enter the keywords "information-security training video" into a

search engine such as Google. After the video film, the team can present its findings from the preliminary evaluation. The immediate goal is to constitute an *information protection working group* to set priorities, determine an action plan, define a timetable and milestones, and formulate policies and procedures to protect corporate information resources. The presenters should name the people you want to see in your working group; all of these people should be contacted before the meeting to be sure that they have agreed in advance to participate in the working group.

The presenters should provide estimates of the time involved and the costs of in-house, and consulting, services, and software. To end the briefing, it is useful to offer upper managers a range of background reading about security. Some managers may be intrigued by this field; the more they learn, the more they will support security efforts. One of the best resources for such sensitization is "Managing Risk from Information Systems: An Organizational Approach."[12] This 67-page summary provides "guidelines for managing risk to organizational operations and assets, individuals, other organizations, and the nation resulting from the operation and use of information systems." The authors state: "The guidelines provided in this special publication have been broadly developed from a technical perspective to be generally useful across a wide range of organizations employing information systems to implement mission and business processes." They also provide extensive cross-indexing to other public documents issued by the National Institute of Standards and Technology (NIST), all of which are freely available online.[13]

66.5 PHASE 3: NEEDS ANALYSIS. The information protection working group should include representatives from every sector of the enterprise. As the group investigates security requirements, the participants' wide experience and perspective will be crucial in deciding which areas to protect most strongly. More important, their involvement is a concrete expression of corporate commitment to a fundamental attitude change in the corporate culture: Security is to be an integral part of the corporate mission.

For example, in a manufacturing firm, the team would include managers and staff from the factory floor, the unions, engineering, equipment maintenance, shipping and receiving, facilities management (including those responsible for physical security), administrative support, sales, marketing, accounting, personnel, the legal department, and information systems. Each of these members of the working group will help improve enterprise security.

If the organization is very large, the group may have to set up subcommittees to deal with specific sectors. Each subcommittee evaluates to what degree the systems and networks are vulnerable to breaches of security. For example, one group could focus on local and campus communications, another on wide area enterprise networks, and a third on electronic data interchange with clients and suppliers.

A typical audit covers the facilities, personnel policies, existing security, application systems, and legal responsibility to stakeholders (owners, shareholders, employees, clients, and the surrounding community). Based on the findings, the subcommittees formulate proposals for improving security. This is where the specialized knowledge obtained from information security specialists and information-security courses will prove especially useful.[14]

66.6 PHASE 4: POLICIES AND PROCEDURES. Once the information protection working group has built a solid floor of understanding of enterprise information security needs, the members are ready to construct the policies and procedures that

meet those needs. The process should start from existing templates and normally takes weeks to months to complete a workable draft.[15]

Genuine participation by all the representatives from every sector of the enterprise is a critical element of success; without a thoroughgoing sense of ownership of the policies, working group members will fail to internalize the new policies. All the members of the working group must become enthusiasts for their collective efforts; in some sense, these people become missionaries engaged in the long-term conversion efforts of phase 5, the implementation of the policies.

66.7 PHASE 5: IMPLEMENTATION. Once the working group members have defined the new or improved security policies, they are about halfway to their goal. The hardest part is ahead: explaining the need for security and the value of the new policies to fellow employees and convincing them to change. Even if they agree intellectually, there is a good chance that their ingrained social habits will override the new rules for at least months and possibly years. The challenge is to overcome these habits.

Chapter 50 shows in detail how to use the insights of social psychology to change corporate culture by working on beliefs, attitudes, and behavior. In addition to the suggestions in that chapter, the information protection working group should organize and deliver awareness and training sessions for all levels of the enterprise:

- Upper management
- Technical support
- Lower-level staff
- Other technical staff

The next sections offer some simple agendas for such preliminary sessions.

66.7.1 Upper Management. Security policies and procedures require management support and sanctions. The transformation of corporate culture should begin at the top. Although it is difficult to coordinate the presence of top executives, the working group should try to organize a half-day executive briefing session on enterprise security. In practice, the group may be able to convince upper management to attend for one or two hours. The focus should be intensely practical, and should show executives how to protect themselves and the enterprise against common dangers. Suggested topics:

- **A review of the business case for improving security.** Industrial espionage, natural and manmade disasters, vandalism
- **Network security.** Protection against eavesdropping and tampering
- **Access controls.** Tokens, biometrics, passwords
- **Encryption.** Email, laptops, provision for emergency data recovery
- **Backup policies** for PCs, networks, and mainframes
- **Security agreements.** Summaries of the policies and procedures to be read and signed annually
- **Need for total support** to convince other staff to comply with security policies

66.7.2 Technical Support. The next target is the technical support group, the people who help explain security policies to users. In a one-day training session, the presentations can cover:

- Everything covered in the executive briefing
- Operating system security provisions
- Security software features
- Changes in operations to comply with new procedures

66.7.3 Lower-Level Staff. Lower-level staff need a half-day session that answers these questions in terms that apply directly to their own work:

- Why should I care about information security?
- What are my obligations as an employee?
- How do I protect the PC I am responsible for against viruses?
- How do I back up my data?
- How do I manage my passwords?
- What must I do if I see someone violating our security policies?

The class ends with participants signing the security agreement.

66.7.4 Other Technical Staff. More intensive training and education are needed for technical staff, such as members of the software development, operations, and network administration groups. More in-depth, specific material will have to be incorporated into their training; however, such training can be spread over a longer time than that for the groups already discussed, because of the rhythm of work and the crucial importance of technical competence for implementation of the policies. Most enterprises rely on outside trainers, specialized off-site or online courses, and certification programs to raise their staff to the appropriate levels of competence.

66.8 PHASE 6: MAINTENANCE. Once the enterprise has begun to integrate a concern for security into every aspect of its work, the issue must be kept fresh and interesting. As described in Chapter 49, successful security awareness programs include amusing posters, interesting videos, occasional seminars on stimulating security topics such as recent frauds or computer crimes, and regular newsletters with up-to-date information. Finally, every employee should regularly reread and sign the annual security agreement. This practice ensures that no one can argue that the organization's commitment to security is a superficial charade.

66.9 CONCLUDING REMARKS. For a secure installation, three things are essential:

1. Sound policies that have been developed with the cooperation of everyone concerned and that are updated periodically
2. Widespread dissemination of those policies, with ongoing observation of their implementation and with frequent training and continual reinforcement
3. Commitment to security on the part of everyone, from top management on down

When these essential elements are in place, the entire organization will function at a more productive level, one at which the possibilities of disruption and damage will have been reduced to a minimum. Nothing less is acceptable.

66.10 FURTHER READING

Barman, S. *Writing Information Security Policies*. New Riders Publishing, 2001.

Brotby, K. *Information Security Governance: A Practical Development and Implementation Approach.* Wiley, 2009.

Calder, A., and S. Watkins. *IT Governance: An International Guide to Data Security and ISO27001/ISO27002*, 5th ed. Kogan Page, 2012.

Fugini, M. G., and C. Bellettini. *Information Security Policies and Actions in Modern Integrated Systems*. IGI Global, 2004.

Greene, S. *Security Policies and Procedures: Principles and Practices*, 2nd ed. Pearson IT Certification, 2013.

Howard, P. D. *FISMA Principles and Best Practices: Beyond Compliance*. Auerbach Publications, 2011.

Jacobs, S. *Engineering Information Security: The Application of Systems Engineering Concepts to Achieve Information Assurance*. Wiley-IEEE Press, 2011.

Johnson, R., and M. Merkow. *Security Policies and Implementation Issues*. Jones & Bartlett Learning, 2010.

Peltier, T. R. *Information Security Policies and Procedures: A Practitioner's Reference*, 2nd ed. Auerbach Publications, 2004.

Smith, R. *The Definitive Guide to Writing Effective Information Security Policies and Procedures*. CreateSpace Independent Publishing Platform, 2010.

Williams, B. L. *Information Security Policy Development for Compliance: ISO/IEC 27001, NIST SP 800-53, HIPAA Standard, PCI DSS V2.0, and AUP V5.0*. Auerbach Publications, 2013.

Wood, C. C. *Information Security Policies Made Easy*, 12th ed. Information Shield, 2012.

Wood, C. C. *Information Security Roles & Responsibilities Made Easy*, 2nd ed. Information Shield, 2005.

66.11 NOTES

1. S. Berinato, "The Fifth Annual Global State of Information Security." *CIO*, August 27, 2007. www.cio.com/article/133600/The_Fifth_Annual_Global_State_of_Information_Security
2. See Chapter 67 in this *Handbook.*
3. For additional ideas in framing questions about security awareness, see Chapter 49. For ideas on appropriate questions dealing with employment practices and policies, see Chapter 45.
4. For more detail about physical security and additional ideas on appropriate questions, see Chapters 22 and 23.
5. For much more information suitable for devising questions about development security, see Chapters 38 and 39.
6. For more information suitable for devising questions about operations security, see Chapter 47.
7. For additional ideas on looking at identification and authentication, see Chapters 28 and 29.

8. See Chapter 41 for more ideas on checking for appropriate levels of antimalware precautions.
9. For additional suggestions to help in framing questions about data backups, see Chapter 57.
10. See Chapter 56.
11. For more ideas on questions that are appropriate in quickly evaluating the state of business resumption planning and disaster recovery, see Chapters 58 and 59.
12. R. Ross, S. Katzke, A. Johnson, M. Swanson, and G. Stoneburner, "Managing Risk from Information Systems: An Organizational Perspective," National Institute of Standards and Technology NIST SP 800-39, Second Public Draft (April 2008), http://csrc.nist.gov/publications/drafts/800-39/SP800-39-spd-sz.pdf
13. The *Further Reading* sections of chapters in this *Handbook* (as well as sections of the chapters themselves) provide a wealth of material for management sensitization.
14. Chapters 54, 63, and 62 have information that will help the information protection working group develop an evaluation plan.
15. Chapter 44 contains many practical suggestions and resources for the content and style of security policies.

CHAPTER 67

DEVELOPING CLASSIFICATION POLICIES FOR DATA

Karthik Raman, Kevin Beets, and M. E. Kabay

67.1 INTRODUCTION 67·1

67.2 WHY PERFORM DATA CLASSIFICATION? 67·2

67.3 DATA CLASSIFICATION'S ROLE IN INFORMATION SECURITY 67·2

67.4 LEGAL REQUIREMENTS, COMPLIANCE STANDARDS, AND DATA CLASSIFICATION 67·3

- 67.4.1 Privacy Act of 1974 67·3
- 67.4.2 Family Educational Rights and Privacy Act (FERPA) 67·3
- 67.4.3 Health Insurance Portability and Accountability Act (HIPAA) 67·4
- 67.4.4 Gramm-Leach-Bliley Act 67·4
- 67.4.5 Sarbanes-Oxley Act 67·5
- 67.4.6 The Federal Rules of Civil Procedure 67·5
- 67.4.7 Compliance Standards 67·6

67.5 DESIGNING AND IMPLEMENTING DC 67·6

67.6 DATA CLASSIFICATION SOLUTIONS 67·6

- 67.6.1 General Observations 67·7
- 67.6.2 Varonis: Five Steps to Faster Data Classification 67·7
- 67.6.3 Microsoft's Data Classification Toolkit 67·8

67.7 EXAMPLES OF DATA CLASSIFICATION SCHEMAS 67·8

- 67.7.1 Universities 67·9
- 67.7.2 U.S. Federal Government 67·9
- 67.7.3 State Governments 67·12

67.8 CONCLUDING REMARKS 67·15

67.9 FURTHER READING 67·16

67.10 NOTES 67·17

67.1 INTRODUCTION. A figure appears from the bushes on a dark and stormy night. Silently, this figure slips past two crisscrossing guards marching in lock-step. Once past the guards and inside the building, a flashlight flickers to life and begins a slow dance around a cluttered office. The beam freezes. It illuminates an envelope that is stamped with large red letters:

TOP SECRET

The *Top Secret* label is likely the most recognized part of an oft-used schema for specifying the confidentiality requirements of data. This labeling is known as *data classification*. Data classification (DC) characterizes data so that its custodians can decide how to organize, view, edit, value, share, use, and protect data in compliance with security requirements.

Historically, data classification has been used by the government and military. Today however, DC has increasingly become a necessity for businesses as well.[1] The use of computers has only served to escalate this necessity. With data becoming available more readily and going to and from many widespread sources, it is imperative to be able to understand the aspects of these data so that they can be stored and controlled easily and securely.

This chapter highlights why DC is necessary, how it relates to information security, its design and implementation in an enterprise, hardware and software solutions that can assist in performing it, and some points to consider when implementing it.

67.2 WHY PERFORM DATA CLASSIFICATION? Information life cycle management (ILM) is a combination of processes and technology that allow for the control of data throughout its life cycle.[2] This procedure ensures that information technology (IT) professionals can manage data from the point of its inception until the time the data has outlived its usefulness. DC is an important part of ILM.

For a real-world example, consider an audit on a business as part of a HIPAA-compliance check. If a data administrator has not spent the time to classify data, not only will the audit take longer than it should, but the company may be fined for lacking controls on this data simply because they do not understand what their data consists of. An example of such a sticky violation would be the posting of patient records by a clinic on a public-facing Website. If data custodians have not classified data types, how can they be expected to know what to protect and how to disseminate it securely and reliably?

Although legal requirements and compliance fines are a major reason that DC is performed, they are by no means the only motivators for a business to carry out DC. A business stands to gain increased productivity and cost savings using DC. Allowing data to flow freely yet securely is critical to business objectives and success.

One proposed solution to the problem of growing data volumes (and the energy required to run the systems that store the data) has been to reduce the amount of data being accessed on these hosts by purging the unnecessary and using only the requisite data. ILM including DC are essential to determine if a particular set of data has outlived its usefulness.

In summary, DC can bring an organization the following benefits:

- Compliance with data standards and legal requirements
- Streamlined and secure sharing of data
- More efficient data storage and retrieval
- Tracking of data through ILM

67.3 DATA CLASSIFICATION'S ROLE IN INFORMATION SECURITY. Information security, including data classification, must consider the value of data. In addition to the operational value of data for the organization depending on it, data may be valuable to competitors or rapacious nation states with a policy of systematic industrial espionage. Disgruntled and dishonest employees may be able to sell it to a

competitor. Data regarding an organization's clients and other data subjects, including personally identifiable information (PII) (e.g., credit card data) must be safeguarded against loss or theft; failure may result in legal action against that company. Data loss prevention (DLP) strategies and compliance hinge on accurate DC policies and DC implementation.

Guidelines such as those from the Federal Financial Institutions Examination Council (FFIEC) highlight the need to secure data as an integral part of DC:

> Classifying data allows the institution to ensure consistent protection of information and other critical data throughout the system. Classifying systems allows the institution to focus its controls and efforts in an efficient and structured manner. Systems that store or transmit data of different sensitivities should be classified as if all data were at the highest sensitivity. Classification should be based on a weighted composite of all relevant attributes.[3]

DC is an essential means to assist the data custodian to determine the type and value of specific data categories. DC thus helps determine how data will be protected during transport and in storage.

67.4 LEGAL REQUIREMENTS, COMPLIANCE STANDARDS, AND DATA CLASSIFICATION.

Legal requirements and compliance standards may make the marriage of DC and security a *mandatory* step for some businesses. A brief overview of such standards and compliance codes follows.

The amount of data that passes through an average-sized enterprise's computers and networks daily can run into terabytes. Regulations and compliance standards may make it necessary that the enterprise's IT staff be able to manage all of this data. Whatever the regulation or standard, an enterprise will need to be able search through all its data, no matter where it lies in the enterprise's computers, networks, or storage, to discover, preserve, and produce the data to meet a legal request or compliance requirement.

The use of DC to meet legal requests and for compliance has been hastened by the advent of the paperless office. Just 20 years ago, an enterprise could reasonably manage its data for compliance if its (physical) filing system were well organized. In today's digital office, however, an enterprise's data can be scattered over any number of hosts, in several geographic locations, across Storage Area Networks ensconced in offsite data centers, or any combination of each of these. Without a concerted enterprise-wide automation and DC implementation, meeting regulatory requirements will be nearly impossible.

Compliance to a number of laws and standards is facilitated by DC. Some of the key laws and standards requiring data classification are enumerated below.

67.4.1 Privacy Act of 1974.

The Privacy Act of 1974 aims to regulate the disclosure of records by government agencies. The Computer Matching and Privacy Protection Act of 1988 was grafted into the Privacy Act of 1974 and introduced provisions for "agencies to follow when engaging in computer-matching activities; provide matching subjects with opportunities to receive notice and to refute adverse information before having a benefit denied or terminated; and require that agencies engaged in matching activities establish Data Protection Boards to oversee those activities."[4]

67.4.2 Family Educational Rights and Privacy Act (FERPA).

Educational institutions that receive U.S. Government funding have to contend with the requirements of FERPA, which regulates the handling of educational data. There are

many sections of FERPA on information disclosure that are relevant to DC. Here are some examples[5]:

- 99.12 What limitations exist on the right to inspect and review records?
- 99.33 What limitations apply to the re-disclosure of information?
- 99.34 What conditions apply to disclosure of information to other educational agencies or institutions?
- 99.35 What conditions apply to disclosure of information for Federal or State program purposes?
- 99.36 What conditions apply to disclosure of information in health and safety emergencies?
- 99.37 What conditions apply to disclosing directory information?

67.4.3 Health Insurance Portability and Accountability Act (HIPAA). One of the most important regulations that medical organizations must consider is the Health Insurance Portability and Accountability Act (HIPAA) of 1996. Title II of HIPAA, "Preventing Health Care Fraud and Abuse; Administrative Simplification; Medical Liability Reform," mandates a set of standards for the management of health information. The contents of Part C of Title II, "Administrative Simplification" address security[6]:

- Sec. 1171. Definitions.
- Sec. 1172. General requirements for adoption of standards.
- Sec. 1173. Standards for information transactions and data elements.
- Sec. 1174. Timetables for adoption of standards.
- Sec. 1175. Requirements.
- Sec. 1176. General penalty for failure to comply with requirements and standards.
- Sec. 1177. Wrongful disclosure of individually identifiable health information.
- Sec. 1178. Effect on State law.
- Sec. 1179. Processing payment transactions.

67.4.4 Gramm-Leach-Bliley Act. The Gramm-Leach-Bliley Act (GLBA) of 1999 introduced reform for the financial sector. Title V of GLBA imposes privacy requirements on the handling of information by financial organizations: "It is the policy of the Congress that each financial institution has an affirmative and continuing obligation to respect the privacy of its customers and to protect the security and confidentiality of those customers' nonpublic personal information." The contents of "Title V—Privacy" are as follows[7]:

- Subtitle A—Disclosure of Nonpublic Personal Information
  - Sec. 501. Protection of nonpublic personal information.
  - Sec. 502. Obligations with respect to disclosures of personal information.
  - Sec. 503. Disclosure of institution privacy policy.

- Sec. 504. Rulemaking.
- Sec. 505. Enforcement.
- Sec. 506. Protection of Fair Credit Reporting Act.
- Sec. 507. Relation to State laws.
- Sec. 508. Study of information sharing among financial affiliates.
- Sec. 509. Definitions.
- Sec. 510. Effective date.

- Subtitle B—Fraudulent Access to Financial Information
  - Sec. 521. Privacy protection for customer information of financial institutions.
  - Sec. 522. Administrative enforcement.
  - Sec. 523. Criminal penalty.
  - Sec. 524. Relation to State laws.
  - Sec. 525. Agency guidance.
  - Sec. 567. Reports.
  - Sec. 527. Definitions.

67.4.5 Sarbanes-Oxley Act. The Sarbanes-Oxley (SOX) Act was introduced in 2002 in the wake of a slew of accounting scandals involving a few large U.S. corporations. The scope of SOX is broad and many of its sections may have relevance for DC. For example, section 404, "Management assessment of internal controls," makes it the responsibility of management for "establishing and maintaining an adequate internal control structures and procedures for financial reporting."[8]

67.4.6 The Federal Rules of Civil Procedure. The Federal Rules of Civil Procedure (FRCP) are a set of rules for the trial of noncriminal cases in federal courts in the United States. The most recent editions of the FRCP include sections that address electronically stored information (ESI).[9]

Rule 26, "General Provisions Governing Discovery; Duty of Disclosure," states that a party should be able to provide[10]:

> … a copy of, or a description by category and location of, all documents, electronically stored information, and tangible things that are in the possession, custody, or control of the party and that the disclosing party may use to support its claims or defenses, unless solely for impeachment; …

Rule 34, "Production of Documents, Electronically Stored Information, and Things and Entry Upon Land for Inspection and Other Purposes," states:

> … Any party may serve on any other party a request (1) to produce and permit the party making the request, or someone acting on the requestor's behalf, to inspect, copy, test, or sample any designated documents or electronically stored information …

Further, Rule 34 states, "… a party who produces documents for inspection shall produce them as they are kept in the usual course of business or shall organize and label them to correspond with the categories in the request …"

67.4.7 Compliance Standards. Many of the standards discussed in various chapters in this *Handbook* explicitly support DC. Examples include:

- ISO/IEC 27000 family of standards
- Control Objectives for Information and related Technology (COBIT)
- Federal Financial Institutions Examination Council (FFIEC) regulations for finances
- Department of Defense (DoD) issuances for defense contracting
- Media, Telecom: Federal Communications Commission (FCC) regulations
- Food and Drug Administration (FDA) regulations for the life sciences

67.5 DESIGNING AND IMPLEMENTING DC. In this section we outline how a business can design and implement DC. Both tasks must be carried out by IT or INFOSEC staff with the support and collaboration of management. These tasks should be iterative.

The process to design DC can be modeled as follows:

- Obtain management approval.
- Study the organization's business continuity plan (BCP) and note the organization's current IT assets and storage-management processes.
- Present the benefits of DC to the heads of business units (BUs).
- Survey users in various BUs about how they store, retrieve, and edit data and how they would like to see their data organized and labeled.
- List the revenue-generating and mission-critical usage of each BU's data. Understand how each unit uses and manipulates data specific to their needs, and how those data may be shared across other units.

Once the above steps are undertaken, the DC team can devise a data-labeling scheme that takes into account all BUs and their interactions.

The process of implementation of DC can be modeled as follows:

- Obtain management approval
- Map data-labeling scheme to available hardware, networks, systems, and storage
- Apply automation or DC tools where relevant
- Guide users through adoption of new DC scheme and solicit feedback

Finally, the DC team should report the results of DC design and implementation to management. A refined data classification report would include a service-level agreement (SLA) for data usage as well as a comprehensive model of costs.[11] One of the major costs to announce would be for new hardware, networks, storage, or software needed to implement DC.

67.6 DATA CLASSIFICATION SOLUTIONS. Software solutions for DC have been driven mostly by the advances in data-storage optimization. Some

examples of the storage advances that helped to spark these solutions include the following:

- Virtualization—the ability to logically organize multiple physical locations
- De-duplication—technology to reduce duplicate data
- Serial Advanced Technology Attachment (SATA)-based disk arrays—cost-effective storage media

67.6.1 General Observations. Knowing the type of data being managed is imperative when deciding how to secure it, store it, and manage its information life cycle. This process can be involved and extremely time consuming.

There are several vendors who provide software solutions to assist in this classification process. Most of these solutions run on rack-mounted appliance. In fact, an entire market has sprung up—information and classification management (ICM)—to provide tools and consultation for DC.

Some of the features that a data custodian may want to consider when making a software solution decision include:

- Policy-based data-type discovery—DC based on the policies outlined by the data custodian. These policies can later be used to automatically enforce a business's storage policies
- File metadata classification—DC using a file's metadata and file content
- Multiple file system management—seamless classification across multiple file system types
- Compliance and legal considerations—DC that includes usage of legal metadata, retains file ownership chains, and uses pattern-matching for sensitive data such as Social Security numbers
- Report style—DC based on the utility of the reports on data that has been discovered, labeled, and classified

67.6.2 Varonis: Five Steps to Faster Data Classification. In a white paper, Varonis Systems, makers of the Varonis Data Governance suite, summarize their proposed five-step DC approach as follows:

1. **Identify Data Owners:** Use Varonis DatAdvantage to identify and assign data owners for important data without a known owner. In this example, we'll say "Dave Smith" was determined to be the owner for the SharePoint site "Project X Resources."
2. **Define Data of Interest:** Work with Dave Smith to define what he considers sensitive. In this example, Dave says data containing the phrase "Project Budget" or the word "Secret" is sensitive.
3. **Use Meta-Data to Focus and Accelerate**: Dave says the sensitive data on his site should not be accessible to the Marketing or Operations teams. Use Varonis IDU Data Classification Framework to find data on the Project X Resources SharePoint site that is accessible by the Marketing or Operations teams, and contains sensitive content.

4. **Report and Remediate**: Dave receives reports about sensitive data that is accessible by the Marketing and Operations team and works with IT Staff to change data permissions.
5. **Rescan Data**: The Project X Resources site is rescanned weekly, and Dave gets reports showing whether any newly added or modified data violates his specified access policies.[12]

67.6.3 Microsoft's Data Classification Toolkit. Microsoft makes available a free "Data Classification Toolkit,"[13] which it introduces as follows:

> The Data Classification Toolkit for Windows Server 2012 is designed to help organizations:
>
> - Identify, classify, and protect data on file servers in private and public clouds.
> - Take advantage of new features and technologies in Windows Server 2012, as well as support hybrid environments with file servers running Windows Server 2012 and Windows Server 2008 R2 SP1.
> - Easily configure default central access policy across multiple servers.
> - Build and deploy policies cost-effectively to protect critical information.
>
> Use the Data Classification Toolkit to help your organization successfully plan and maintain data classification programs in these critical areas:
>
> - Identifying applicable IT GRC authority documents.
> - Defining corresponding classification policies.
> - Preserving evidence that demonstrates the implementation of effective controls.
>
> When used with the IT GRC Process Management Pack SP1 for System Center Service Manager, organizations can easily map classification requirements from various authority documents, such as PCI DSS and NIST 800-53. In addition, your organization can build and deploy its data classification policies, implement controls through Windows Server 2012 File Classification Infrastructure, and provide evidence of data classification policies to auditors. . . .
>
> Streamline your compliance experience with new features in the Data Classification Toolkit that will help you identify, classify, and protect data in private and public clouds. The toolkit supports file servers running Windows Server 2012 and Windows Server 2008 R2 SP1. In addition to configuring File Classification Infrastructure (FCI), the latest version of the toolkit allows you to manage Central Access Policy across file servers. The toolkit enhances the user experience by providing scenario-based wizards that you can use to configure, export, import, and compare file classifications, as well as manage central access policy on your file servers. It provides tools to provision user and device claim values and central access policy across a forest to help simplify configuring Dynamic Access Control in Windows Server 2012. The toolkit also provides a new report template that you can use to review existing central access policy on file shares. . . .[14]
>
> The Microsoft Data Classification Toolkit (the "software") is intended to help organizations simplify their ability to search, identify, and apply rules to data they specify. The software provides sample search expressions and rules that can be used to assist with compliance activities conducted by your organization's IT professionals, auditors, accountants, attorneys, and other compliance professionals. The software does not replace those professionals. . . .[15]

67.7 EXAMPLES OF DATA CLASSIFICATION SCHEMAS. Security policy and practice can vary widely in design and implementation across different organizations. So too can DC design and implementations.

67.7.1 Universities. For example, here are some differences in the implementation of DC in four U.S. universities:

- The George Washington University's Data Classification Security Policy states, "Data is a critical asset of the University. All members of the University community have a responsibility to protect the confidentiality, integrity, and availability of data generated, accessed, modified, transmitted, stored, or used by the University, irrespective of the medium on which the data resides and regardless of format (such as in electronic, paper, or other physical form)." Three categories of data are defined, "Public," "Official Use Only," and "Confidential." The policy gives examples of each category.[16]
- Stanford University's Data Classification Guidelines Web page states the case for the classification of the University's data clearly and concisely. It defines three categories of data, "Public," "Sensitive," and "Restricted," and tabulates some legal requirements, reputation risks, other risks, access restrictions, and examples against these categories.[17]
- The University of Missouri's DC policy outlines the case for DC and defines four categories of data, "Public," "Confidential," "Restricted," and "National Security Interest." It lists three network zones, each with its own security requirements. It tabulates various set up and usage requirements for each data category.[18]
- Angelo State University's DC policy asks for three DC policies, provides examples of how data in each category can be lost, and sketches what the impact of losses could be.[19]

For each of the above examples, we can see that the DC policies and procedures could vary based solely on the labels chosen. Those entities that use fewer labels will most likely have fewer policies and procedures across a broader spectrum of data as opposed to those who choose to divide their data into more categories. For example, if a business chooses to label all data as "Classified," a single policy would be in place that all users would need to follow, as opposed to a business that chose to label data with many different labels. In practice, the number of labels or categories would most likely be somewhere between three and six.

67.7.2 U.S. Federal Government. The Federal Information Processing Standards Publication 199, "Standards for Security Categorization of Federal Information and Information Systems," provides the baseline for all U.S. Federal-Government agencies in their DC programs.[20] The standard is based on the outmoded C-I-A model (Confidentiality, Integrity, Availability)[21] and weighs the potential impact as low, moderate, or high. These levels are defined as follows:

> The potential impact is LOW if—
>
> —The loss of confidentiality, integrity, or availability could be expected to have a limited adverse effect on organizational operations, organizational assets, or individuals.
>
> AMPLIFICATION: A limited adverse effect means that, for example, the loss of confidentiality, integrity, or availability might: (i) cause a degradation in mission capability to an extent and duration that the organization is able to perform its primary functions, but the effectiveness of the functions is noticeably reduced; (ii) result in minor damage to organizational assets; (iii) result in minor financial loss; or (iv) result in minor harm to individuals.

The potential impact is MODERATE if—

—The loss of confidentiality, integrity, or availability could be expected to have a serious adverse effect on organizational operations, organizational assets, or individuals.

AMPLIFICATION: A serious adverse effect means that, for example, the loss of confidentiality, integrity, or availability might: (i) cause a significant degradation in mission capability to an extent and duration that the organization is able to perform its primary functions, but the effectiveness of the functions is significantly reduced; (ii) result in significant damage to organizational assets; (iii) result in significant financial loss; or (iv) result in significant harm to individuals that does not involve loss of life or serious life threatening injuries.

The potential impact is HIGH if—

—The loss of confidentiality, integrity, or availability could be expected to have a severe or catastrophic adverse effect on organizational operations, organizational assets, or individuals.

AMPLIFICATION: A severe or catastrophic adverse effect means that, for example, the loss of confidentiality, integrity, or availability might: (i) cause a severe degradation in or loss of mission capability to an extent and duration that the organization is not able to perform one or more of its primary functions; (ii) result in major damage to organizational assets; (iii) result in major financial loss; or (iv) result in severe or catastrophic harm to individuals involving loss of life or serious life threatening injuries.

Federal departments have references to internal policies that explicitly address data classification. Here are a few examples:

- **Agriculture**: "Sensitive Security Information means unclassified information of a sensitive nature, that if publicly disclosed could be expected to have a harmful impact on the security of Federal operations or assets, the public health or safety of the citizens of the United States or its residents, or the nation's long-term economic prosperity… ."[22]
- **Defense (DoD)**: A sample policy template available from the Defense Information Systems Agency has the following "Data Rules" governing the treatment of *sensitive* data from "Large Service Applications (LSAs)." The guidelines explicitly state that the references to LSA may be modified to suit each agency.

- All LSA data/information is highly sensitive and should be dated and marked as such. External labeling shall include special handling instructions (e.g., log/inventory identifiers, controlled access, special storage instructions, release, or destruction dates).
- All LSA information transported through the mail or courier/messenger service shall be double-sealed, the second envelope shall be marked "CONFIDENTIAL Designated Official Only."
- The receipt and delivery of sensitive data must be monitored and accounted for to ensure that data is not lost and potentially compromised while in transit.
- Sensitive data shall only be given to those employees with a need to know and who have authorized access in the performance of their official duties.
- Sensitive information must not be left unattended, even temporarily. Sensitive data must remain in the employee's physical control at all times. Sensitive material should be kept in a secure safe or a locked cabinet and returned to the safe each evening or during any lunch periods or breaks greater than 30 minutes.
- Sensitive information shall be turned over or be put out of sight when visitors are present.
- Physical, environmental protection controls shall be provided for sensitive data contained in a media storage vault or library.
- Sensitive information shall not be discussed outside LSA restricted areas.
- Sensitive information being hand-carried must be kept with the individual and protected from unauthorized disclosure… .[23]

- **Energy (DoE)**: "Security Framework for Control System Data Classification and Protection … presents a data classification process that gives utility administrators, control engineers, and IT personnel a cohesive approach to deploying efficient and effective process control security."[24] The Executive Summary introduces the classification scheme as follows: "This document presents a data classification process that gives utility administrators, control engineers, and IT personnel a cohesive approach to deploying efficient and effective process control security. The fundamental goal is a clear delineation of control system data that will enable effective implementation of security techniques and technologies so the control system can function as required in the face of threats. Once created, the data classification security framework will help reduce the risk of energy disruptions due to control system failure by securing data critical to the operation of the control system."[25]
- **Health and Human Services (HHS)**: The *Federal Register* included a discussion in the "Health Insurance Reform: Security Standards" report in 2003 in which "One commenter stated that a data classification policy, that is a method of assigning sensitivity ratings to specific pieces of data, should be part of the final regulations." The response from HHS was, "We did not adopt such a policy because this final rule requires a floor of protection of all electronic protected health information. A covered entity has the option to exceed this floor. The sensitivity of information, the risks to and vulnerabilities of electronic protected health information and the means that should be employed to protect it are business determinations and decisions to be made by each covered entity."[26]
- **Homeland Security (DHS)**: Some of the relevant language pertaining to data classification includes the following definitions and orders:

  - *Confidential*: A level of classification applied to information, the unauthorized disclosure of which reasonably could be expected to cause damage to the national security that the original classification authority is able to identify or describe.
  - *Secret*: Level of classification applied to information, the unauthorized disclosure of which reasonably could be expected to cause serious damage to the national security that the original classification authority is able to identify or describe.
  - *Top Secret*: Level of classification applied to information, the unauthorized disclosure of which reasonably could be expected to cause exceptionally grave damage to the national security that the original classification authority is able to identify or describe.
  - *Top Secret Document Register*: A TOP SECRET Document Register will be maintained by the TSCO to record the receipt, disposition, and destruction of TOP SECRET information. DHS Form 11000-03, Document Control Register, Top Secret National Security Information, may be used for this purpose. Other forms or automated registers created by an organizational element may also be used, but, at a minimum, must contain the information included in DHS Form 11000-03. Electronic registers must be backed up after each transaction to ensure no loss of accountability.
  - *Secret and Confidential*: Except as required by the originator or as specified for certain categories of SECRET and CONFIDENTIAL information, there is no requirement to maintain accountability records or conduct inventories for SECRET and/or CONFIDENTIAL information. However, heads of organizational elements may mandate the use of accountability records within their respective elements at their discretion.[27]

- **State**: In the Department of State's "IT Strategic Plan: Fiscal Years 2006–2010," the authors discussed changes in the data-classification policies and access to different classification levels as follows:

The following list exemplifies the types of decisions that must be made now and in the future:

- Wireless laptops and networks, providing access from home and on travel
- Integration of voice mail and e-mail
- Cell phone access to e-mail
- Re-examination of data classification policies and practices to ensure that information is as accessible as possible
- Classified Instant Messaging
- PDA access to unclassified networks for e-mail and document browsing
- Laptop access to classified and unclassified networks
- PDA access to classified networks
- Inclusion of Sensitive But Unclassified (SBU)/NOFORN on OpenNet
- Inclusion of EXDIS on ClassNet
- Protection of Personal Identifiable Information; particularly on laptops and other remote devices[28]

67.7.3 State Governments. The following shows extracts from a few state government documents on DC.

New York:

This document includes both the Information Classification and Control Policy and the Information Classification and Control Standard. The Policy is part of the broader Information Security Policy (P03-002) and is included in this document for ease of reference.

In order to facilitate the process of classifying information assets an Information Classification Manual, Information Asset Classification Worksheet, Information Control Charts and Glossary of Information Security Controls are provided in the Appendices. The Information Classification Manual, in conjunction with the Worksheet, provides a process for classifying information assets and contains the minimum mandatory questions that must be answered when classifying information. The Control Charts contain the mandatory baseline controls that must be implemented based on the information classification. The Glossary contains explanations for each control. A SE [State Entity] may add more questions and/or controls but may not alter or remove the original questions and controls.

Please note the Worksheet and the Control Charts and Glossary are available in an automated tool called the Information Asset Classification System (IACS). This application is available to all New York State governmental entities utilizing NYS Directory Services… .

The classification of information pursuant to this Policy and application of appropriate controls to that information do not alter the responsibility of the SE to comply with the records retention and disposition requirements of the Arts and Cultural Affairs Law or its responsibility to make records available for public inspection and copying under the provisions of the Freedom of Information Law. The process of classifying information pursuant to this Policy may, however, serve as a basis for a SE to evaluate the retention and disposition schedules currently in effect for its records and, where appropriate, consider revising those schedules as a means of managing the records that must be protected by the SE. Similarly, the classification process can facilitate the accurate and efficient application of the exemptions from disclosure enumerated in the Freedom of Information Law by providing a framework for the comprehensive assessment of the SE's information assets.

A. All information assets must have an information owner established within the SE's lines of business. The information owner will be responsible for assigning the information classification, determining access privileges of users or groups of users based on job duties, and overseeing daily decisions regarding information asset management. Periodic reviews will be performed by the information owner to confirm the classification of, or reclassify, the information asset.

B. Each classification will have an approved set or range of controls. If SE information is stored by a third-party, the information owner's SE is responsible for communicating requirements of this Policy and Standard to the third-party and addressing them in third-party agreements as they relate to the SE's data.

C. An information asset must be classified based on the highest level necessitated by its individual data elements.

D. All Personal, Private, or Sensitive Information (PPSI) shall be classified with a confidentiality of high.

E. Merging of information which creates a new information asset or situations that create the potential for merging (e.g., backup tape with multiple files) must be evaluated to determine if a new classification of the merged data is warranted.

F. If the SE is unable to determine the confidentiality classification of information stored on electronic storage media, the information must be assumed to have a high confidentiality classification and, therefore, is subject to high confidentiality controls.

G. All reproductions of information in its entirety must carry the same confidentiality classification as the original. Partial reproductions need to be evaluated to determine if a new classification is warranted.

H. A written or electronic inventory of all SE information assets must be maintained.[29]

Ohio:

This state policy provides a high-level data classification methodology to state agencies for the purpose of understanding and managing data and information assets with regard to their level of confidentiality and criticality. Accurate identification provides a basis to employ an appropriate level of security… .

State agencies shall establish a data classification policy in compliance with this state policy. Each agency shall serve as a classification authority for the data and information that it collects or maintains in satisfaction of its mission… .

Data Classification Labels. The classification of data is a critical tool in defining and implementing the correct level of protection for state information assets. Such classifications are a prerequisite to establishing agency guidelines and system requirements for the secure generation, collection, access, storage, maintenance, transmission, archiving and disposal of state data. Data classification shall be part of the agency's overall risk assessment process as outlined in Ohio IT Policy ITPB. 1, "Information Security Framework."

Agencies shall label data for both confidentiality and criticality. Such classification labels are defined at a high level and represent broad categories of information. State and federal law may also require specific labels.

Confidentiality. The confidentiality label identifies how sensitive the data is with regard to unauthorized disclosure. Data shall be assigned one of three labels for confidentiality:

- **Public**. The "public" label includes information that must be released under Ohio public records law or instances where an agency unconditionally waives an exception to the public records law.
- **Limited Access**. The "limited-access" label applies to information that an agency may release if it chooses to waive an exception to the public records law and places conditions or limitations on such a release.
- **Restricted**. The "restricted" label applies to information, the release of which is prohibited by state or federal law. This label also applies to records that an agency has discretion to release under public records law exceptions but has chosen to treat the information as highly confidential.

Criticality. The criticality label identifies the degree of need for data to maintain its integrity and availability. Data shall be assigned one of four labels for criticality:

- **Low.** The loss of data integrity or availability would result in insignificant or no financial loss, legal liability, public distrust, or harm to public health and welfare.

- **Medium.** The loss of data integrity or availability would result in limited financial loss, legal liability, public distrust, or harm to public health and welfare.
- **High.** The loss of data integrity or availability would result in significant financial loss, legal liability, public distrust, or harm to public health and welfare.
- **Very High.** The loss of data integrity or availability would result in catastrophic financial loss, legal liability, public distrust, or harm to public health and welfare.[30]

Texas:

Data Classification provides a framework for managing data assets based on value and associated risks and for applying the appropriate levels of protection as required by state and federal law as well as proprietary, ethical, operational, and privacy considerations. All [AGENCY] data, whether electronic or printed, should be classified. The data owner, who is responsible for Data Classification, should consult with legal counsel on the classification of data as Confidential, Agency-Sensitive, or Public. Consistent use of data classification reinforces with users the expected level of protection of [AGENCY] data assets in accordance with [AGENCY] security policies.

The purpose of the [AGENCY] Data Classification Policy is to provide a foundation for the development and implementation of necessary security controls to protect information according to its value and/or risk. Security standards, which define these security controls and requirements, may include: document marking/labeling, release procedures, privacy, transmission requirements, printing protection, computer display protections, storage requirements, destruction methods, physical security requirements, access controls, backup requirements, transport procedures, encryption requirements, and incident reporting procedures.

Data shall be classified as follows:

- Confidential.
  - Sensitive data that must be protected from unauthorized disclosure or public release based on state or federal law (e.g. the Texas Public Information Act) and other constitutional, statutory, judicial, and legal agreements.
  - Examples of "Confidential" data may include but are not limited to:
    - Personally Identifiable Information, such as: a name in combination with Social Security Number (SSN) and/or financial account numbers
    - Student Education Records
    - Intellectual Property, such as: Copyrights, Patents, and Trade Secrets
    - Medical Records
- Agency-Sensitive.
  - [optional AGENCY defined category that may be identified as: Agency "Security-Sensitive," "Privileged," or "Protected"].
  - Sensitive data that may be subject to disclosure or release under the Texas Public Information Act, but requires additional levels of protection.
  - Examples of "Agency-Sensitive" data may include but are not limited to:
    - [AGENCY] operational information
    - [AGENCY] personnel records
    - [AGENCY] information security procedures
    - [AGENCY] research
    - [AGENCY] internal communications
    - Public. Information intended or required for public release as described in the Texas Public Information Act[31]

West Virginia:

All State data requires classification, which is mandatory for more confidential and critical classes of data… .

- Level 1—Extremely Sensitive Data
  - Extremely sensitive data is the most sensitive data to integrity and confidentiality risks. Disclosure or corruption of this data could be hazardous to life or health… .

- Access is tightly restricted with the most stringent security controls at the system as well as the user level. Failure to maintain the integrity and confidentiality could have severe financial, health, or safety repercussions. Individuals must adhere to very strict rules in the usage of this data… .
  - Examples of extremely sensitive data may include the following: Contents of State law enforcement investigative records;
  - Child and adult protective services client data.

- Level 2—Very Sensitive Data
  - This data is only made available to authorized users and may be protected by federal and State regulations.
  - Access to very sensitive data is restricted to authenticated and authorized individuals who require access to that information in the course of performing their job duties. These are the data elements removed from responses to information requests for reasons of privacy.
  - Examples of very sensitive data may include the following:
    - Social Security numbers;
    - Credit card numbers;
    - Food assistance programs data;
    - Comprehensive law enforcement data;
    - Foster care data;
    - Health, mental health, acute medical care, and medical data;
    - Social Service or Temporary Assistance data; and
    - Tax information.

- Level 3—Sensitive Data
  - This data is made available through open record requests or other formal or legal processes; it includes the majority of the data contained within State government electronic databases.
  - Direct access is restricted to authenticated and authorized individuals who require access to that information in the course of performing their job duties.
  - Examples of sensitive data may include the following:
    - Most data elements in State personnel records;
    - Driver history records;
    - State/federal contracts data;
    - Employment and training program data;
    - Permits data; and
    - Historical records repository data.

- Level 4—Unrestricted Data
  - This data is characterized as being open, public data with no distribution limitations and to which anonymous access is allowed.
  - This type of information is: (1) actively made publicly available by State government; (2) published and distributed freely, without restriction; and (3) available in the form of physical documents such as brochures, formal statements, press releases, reports, web pages, and bulletin boards accessible with anonymous access.
  - Examples of unrestricted data may include the following:
    - Occupational licensing data excluding social security numbers;
    - Agency public websites; and
    - Statewide policies.[32]

67.8 CONCLUDING REMARKS. This chapter has introduced DC, why it is relevant, how it relates to compliance and legal requirements, how one may design and implement it, and provided some examples of how a handful of entities have implemented their DC labeling schema.

In conclusion, here are some important points to remember when dealing with DC:

- A complete solution will most likely not be achieved by using software solutions alone. Since DC is still evolving, these solutions may require customization to fit what a business must accomplish. Thorough research will be required to find a best-fit solution.[33]

- In order to determine what type of data is used across an entire business, it is imperative to interview users from all segments of the business using surveys.
- Locations of data must be written into a DC policy. Since data can reside in many different locations, a policy must be drafted to account for geography and allow for proper storage and retrieval of these data from varied locations.
- An educational policy for users will need to be created for the "Important" data type (or equivalent) level and above. Creators and users of valuable data types must be educated about the policies and procedures that are put into place for the storage, retrieval, and use of their data.
- A complete DC solution will take some time to implement. Because of the complexity of DC, a complete and well-designed implementation will not happen overnight. A realistic estimate should be provided to management when undertaking a DC project.
- Compliance and legal requirements need to be understood well for a particular business. There are several compliance standards that make a business legally liable for certain data types. It is imperative that one understand which of those standards apply to their data when taking on the DC task. Remember too, that legal requirements may also cover investigations that require this data, so understanding data-handling laws is also important.

67.9 FURTHER READING

Anonymous. "Data Classification Methods and Policy." *ComputerWeekly*, April 2009. www.computerweekly.com/report/Data-classification-methods-and-policy

Anonymous, "Sensitive Data Classification and Protection: Overcoming the Challenges to Classify and Protect Sensitive Data at Federal Government Agencies." SECUREIIT, 2008. www.secureit.com/resources/WP_Data_Class_and_Protect.pdf

Buckley, S. "Data Classification." *Internal Auditor*, March 2011. www.theiia.org/intAuditor/itaudit/2011-articles/data-classification

Etges, R., and K. McNeil. "Understanding Data Classification Based on Business and Security Requirements." *ISACA Journal Online*, August 10, 2006. www.isaca.org/Journal/Past-Issues/2006/Volume-5/Documents/jopdf0605-understanding-data.pdf

Fowler, S. "Information Classification—Who, Why, and How." *GIAC Security Essentials Certification*, February 28, 2003. SANS Reading Room. www.sans.org/reading_room/whitepapers/auditing/information-classification-who_846

Heiser, J. "Data Classification Best Practices: Techniques, Methods and Projects." *SearchSecurity*, March 2009. Requires membership to view. http://searchsecurity.techtarget.com/magazineContent/Data-classification-best-practices-Techniques-methods-and-projects

Libicki, M. C., D. R. Frelinger, B. A. Jackson, B. E. Lachman, and C. Ip. *What Should Be Classified? A Framework with Application to the Global Force Management Data Initiative*. Rand Publishing, 2011.

Murray, W. H. "Classification and Labeling of Data." *Thinking About Security* blog, March 13, 2012. http://whmurray.blogspot.com/2012/03/classification-and-labeling-of-data.html

Reed, B. "Data-classification best practices." *NetworkWorld*, January 18, 2007. www.networkworld.com/news/tech/2007/012207-techupdate-data-classification.html

Vogel, V. "Data Classification Toolkit." In *Information Security Guide: Effective Practices and Solutions for Higher Education.* Internet2, last edited December 22, 2010. https://wiki.internet2.edu/confluence/display/itsg2/Data+Classification+Toolkit

67.10 NOTES

1. Dataglobal, "Dataglobal Announces dg classification 2.0, the Industry's First Unified Data Classification Solution," Press release via PRWeb, June 11, 2012, www.prweb.com/releases/2012/6/prweb9590296.htm
2. S. Duplessie, N. Marrone, and S. Kenniston, "The New Buzzwords: Information Lifecycle Management," *Computerworld*, March 31, 2003, www.computerworld.com/hardwaretopics/storage/story/0,10801,79885,00.html
3. Federal Financial Institutions Examination Council, "Information Security Booklet," July 2006, http://ithandbook.ffiec.gov/it-booklets/information-security.aspx
4. 5 U.S.C. §552a, U.S. Congress, The Privacy Act Of 1974, December 31, 1974, http://www.justice.gov/opcl/privstat.htm
5. 20 U.S.C. §1232g, U.S. Congress, Family Educational Rights and Privacy Act (FERPA), August 21, 1974, www.access.gpo.gov/nara/cfr/waisidx_04/34cfr99_04.html
6. 42 USC §201 *et seq*, U.S. Congress, Health Insurance Portability And Accountability Act of 1996, August 21, 1996, www.cms.hhs.gov/HIPAAGenInfo/Downloads/HIPAALaw.pdf
7. 15 USC §6801 *et seq*, U.S. Congress, Text of the Conference Report of Gramm-Leach-Bliley Bill, November 12, 1999, http://banking.senate.gov/conf/somfinal.htm
8. 15 U.S.C. §7241, 18 U.S.C. §1350, U.S. Congress, Sarbanes-Oxley Act of 2002, July 30, 2002, http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=107_cong_bills&docid=f:h3763enr.tst.pdf
9. Kahn Consulting, Inc., "The Federal Rules of Civil Procedure: Meeting the IT and Legal Challenges of the New E-Discovery Rules," Law.com, May 1, 2007, www.law.com/jsp/legaltechnology/detailWP.jsp?id=1190797378589 (registration required)
10. Committee on the Judiciary, One Hundred Ninth Congress, "Federal Rules of Civil Procedure," December 1, 2006, http://judiciary.house.gov/media/pdfs/printers/109th/31308.pdf
11. K. Langdon and J. Merryman, "Data Classification: Getting Started," TechTarget, July 1, 2005 http://searchstorage.techtarget.com/magazineFeature/0,296894,sid5_gci1258224,00.html (registration required).
12. Varonis, "5 Steps to Faster Data Classification," Varonis Website, January 22, 2010, www.varonis.com/go/resources/#Whitepapers (registration required).
13. R. Smith, "How to Use Microsoft's Data Classification Toolkit," *Biztech*, December 1, 2011, www.biztechmagazine.com/article/2011/12/how-use-microsofts-data-classification-toolkit
14. Microsoft, "Data Classification Toolkit," September 10, 2012, http://technet.microsoft.com/en-us/library/hh204743.aspx
15. Microsoft, "Important Information about the Data Classification Toolkit," September 10, 2012, http://technet.microsoft.com/en-us/library/hh367453.aspx

16. CIO, George Washington University, "Data Classification Security Policy," December 6, 2005, http://my.gwu.edu/files/policies/DataClassificationPolicy.pdf
17. Stanford University, "Classification of Data," November 6, 2007, www.stanford.edu/group/security/securecomputing/dataclass_chart.html
18. Division of Information Technology, University of Missouri, MU Data Classification System, August 3, 2007, http://doit.missouri.edu/security/data-classification/
19. Angelo State University, "Data Classification Standard," March 25, 2012, www.angelo.edu/services/technology/it_policies/dataClassificationStandard.php
20. NIST, "Standards for Security Categorization of Federal Information and Information Systems," FIPS PUB 199, February 2004, http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf
21. This model omits consideration of control, authenticity, and utility as defined in Chapter 3 in this *Handbook*.
22. U.S. Department of Agriculture, "Control and Protection of 'Sensitive Security Information,"' Departmental Regulation 3440-002, January 30, 2003, www.fas.org/sgp/othergov/usda3440-02.html
23. NIST, "Sample Generic Policy and High Level Procedures for Marking, Handling, Processing, Storage, and Disposal of Data," August 2, 2000, http://csrc.nist.gov/groups/SMA/fasp/documents/production_io_controls/marking.doc
24. U. S. Department of Energy, "Security Framework for Control System Data Classification and Protection," Office of Electricity Delivery & Energy Reliability, 2013, http://energy.gov/oe/downloads/security-framework-control-system-data-classification-and-protection
25. B. T. Richardson and J. Michalski, "Security Framework for Control System Data Classification and Protection," Sandia National Laboratories Report SAND2007-3888P, July 2007, http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/21-Security_Framework_for_Data_Class.pdf
26. "45 CFR Parts 160, 162, and 164 Health Insurance Reform: Security Standards; Final Rule." *Federal Register* 68, no. 34 (February 20, 2003):8345, note m, www.hhs.gov/ocr/privacy/hipaa/administrative/securityrule/securityrulepdf.pdf
27. U. S. Department of Homeland Security, "Protection of Classified National Security Information: Accountability, Control, and Storage," Management Directive System MD Number 11045, October 4, 2004, www.dhs.gov/xlibrary/assets/foia/mgmt_directive_11045_protection_of_classified_national_security_information_accountability_control_and_storage.pdf
28. U. S. Department of State, "IT Strategic Plan: Fiscal Years 2006–2010," Information Resource Management, Department of State Publication Number 11455, December 2007, www.state.gov/m/irm/rls/92497.htm
29. T. D. Smith, "Cyber Security Policy & Standard PS08-001: Information Classification and Control," Office of Cyber Security, New York State Division of Homeland Security and Emergency Services, revised February 7, 2012, www.dhses.ny.gov/ocs/resources/documents/PS08-001-V1.2-February-7-2012.pdf
30. State of Ohio, "Data Classification." Statewide IT Policy Number ITP-B.11, March 19, 2008, www.oit.ohio.gov/IGD/policy/pdfs_policy/ITP-B.11.pdf
31. Texas Department of Information Resources, "IS Security Policies: Data Classification" template (Word .DOC file), Chief Information Security Office, March 30,

2010, http://www2.dir.state.tx.us/SiteCollectionDocuments/Security/Policies and Standards/data_classification_policy.doc

32. State of West Virginia, "Policy: Data Classification," Office of Technology Policy WVOT-PO1000, January 6, 2010, www.technology.wv.gov/SiteCollectionDocuments/ITC Policy Guidelines/Data Class PO.pdf
33. B. Reed, "Data-Classification Best Practices," *Network World*, January 18, 2007, www.networkworld.com/news/tech/2007/012207-techupdate-data-classification.html

CHAPTER 68

OUTSOURCING AND SECURITY

Kip Boyle, Michael Buglewicz, and Steven Lovaas

68.1 INTRODUCTION 68·2
- 68.1.1 Definitions 68·3
- 68.1.2 Distinctions 68·3
- 68.1.3 Insourcing 68·4
- 68.1.4 Nearshoring 68·5
- 68.1.5 Offshoring 68·5

68.2 WHY OUTSOURCE? 68·6
- 68.2.1 Effectiveness versus Efficiency 68·7
- 68.2.2 Being Effective 68·7
- 68.2.3 Being Efficient 68·8

68.3 CAN OUTSOURCING FAIL? 68·8
- 68.3.1 Why Does Outsourcing Fail? 68·9
- 68.3.2 Intended and Unintended Consequences of Outsourcing 68·10
- 68.3.3 Universal Nature of Risk 68·11
- 68.3.4 Clarity of Purpose and Intent 68·12
- 68.3.5 Price 68·13
- 68.3.6 Social Culture 68·14
- 68.3.7 International Economics 68·14
- 68.3.8 Political Issues 68·15
- 68.3.9 Environmental Factors 68·16
- 68.3.10 Travel 68·16
- 68.3.11 Labor 68·17
- 68.3.12 Intellectual Property Risks 68·17
- 68.3.13 Additional Risks 68·18

68.4 CONTROLLING THE RISKS 68·19
- 68.4.1 Controls on What? 68·19
- 68.4.2 Controlling Outsourcing Risk 68·19
- 68.4.3 Availability Controls 68·19
- 68.4.4 Utility Controls 68·21
- 68.4.5 Integrity and Authenticity Controls 68·21
- 68.4.6 Confidentiality and Possession Controls 68·22
- 68.4.7 Making the Best of Outsourcing 68·23

68.5 OUTSOURCING SECURITY FUNCTIONS 68·23
- 68.5.1 Who Outsources Security? 68·24
- 68.5.2 Why Do Organizations Outsource Security? 68·24
- 68.5.3 What Are the Risks of Outsourcing Security? 68·26
- 68.5.4 How to Outsource Security Functions 68·27
- 68.5.5 Controlling the Risk of Security Outsourcing 68·30

68.6 CONCLUDING REMARKS 68·30

68.7 FURTHER READING 68·31

68.8 NOTES 68·31

68.1 INTRODUCTION. Since this chapter was first written in 2008, the world has seen significant transformation, driven to a large extent by economic factors, social networking phenomena, the cloud, and the increasingly pervasive need for always-on, always-available information and entertainment.

Global economic leadership is in transition. Europe and the European Union (EU) are locked in a struggle to remain economically viable, and EU leadership is forced into making decisions regarding who can stay in the EU. The BRIC countries (Brazil, Russia, India, and China) are all facing serious economic and political challenges and all the while unemployment in the United States remains at levels not seen since the Great Depression.

Global, regional, and local businesses around the world everyday work in this dynamic environment. Current multiyear contracts with outsourcers were likely executed in circumstances differing from today's rapidly evolving environment.

Outsourcing and security can mean outsourcing certain corporate activities to a third party and doing so in a security-conscious fashion; for example, business A hires business B to be business A's helpdesk. A second interpretation is that business A hires business B to provide corporate technology security.

The term "outsourcing" has come to identify several distinct concepts, each requiring a different risk-management strategy. In this chapter, we examine today's practice of outsourcing and the effects and considerations it has, or should have, on the work of information assurance professionals.

Organizations (companies, nonprofits, government agencies, etc.) outsource to gain efficiencies and effectiveness. The efficiencies gained, however, do have consequences. An outsourcing strategy of *implement-it-and-forget-it* is unwise, as the outsourcing environment can change quickly and dramatically. Within very recent history, outsourcing primarily referred to the movement of tasks from within a corporate environment to another corporation perhaps better suited or more efficient at those tasks. Now the term *outsourcing* also refers to transferring operations and ownership of portions of a corporate infrastructure, including key computing functions, to a cloud vendor, a solutions integrator, or other entity providing computing services. Thus, outsourcing no longer refers to a call center or provider, but instead covers data possession, movement, processing, protection, and retirement all done in a cloud spread across multiple servers in multiple data centers in multiple areas of the world.

Michael Cooney points out some of the significant considerations and problems facing the outsourcing world at end of the first decade of the twenty-first century:

- Security problems
- China's rise
- India's growth
- India's turmoil
- Blunders
- H-1B visa trials
- Small business outsourcing
- Managed services[1]

Despite the negative overtones that sometimes accompany the practice of outsourcing, there can be great benefits. Here are a few:

- A corporation's self-knowledge improves dramatically when it is compelled to articulate, explain, and define its corporate mission to others.

- Outsourcing clarifies and institutionalizes roles, goals, and measures of success.
- One of the largest and most important aspects is that outsourcing done successfully demands risk identification and formalized mitigation activities that might otherwise not occur within the corporation.
- Outsourcing allows a company to focus on corporate core competencies, thus making gains in:
  - Effectiveness (i.e., doing the tasks for which you have the potential to produce the desired results)
  - Efficiencies (i.e., using an optimal amount of time and energy to complete a task)
  - Focus (i.e., concentrating on key business objectives)
  - Discipline (i.e., holding the corporation accountable for corporate objectives, and holding the vendor to clearly and contractually defined objectives)
  - Appropriate utilization of high-value employees performing high-value business functions
- An organization gains or intends to gain financial improvements by moving work to vendors that can perform the work at a lower cost and/or at a greater level of specialization.

The outsourcing of operational computing functions to the cloud or solution providers allows a company to focus less on the tactical production of data for business analysis. Rather, when an organization doesn't need to focus on the tactical operation of keeping farms of data servers and complicated data storage techniques, that organization can focus on consumption and analysis of the data and the maturation of business knowledge and theoretically, business advantage.

68.1.1 Definitions

Vendor or contractor—an arm's-length entity providing an outsourced service

Organization or business—the entity contracting for its products or services with a vendor

Outsourcing—the fulfillment of a specific business function or functions by contracting with a vendor to perform within the vendor's own facilities

Insourcing—the commonplace use of contract or noncompany employees to fulfill certain business functions within the physical and logical corporate boundaries

Cloud services—data center and or software services offered by commercial computing corporation such as email, communications, collaboration technologies, and data services

68.1.2 Distinctions. The risks and considerations of outsourcing an inbound call center, outsourcing a corporate IT function, insourcing a corporate finance function, insourcing an HR function, outsourcing corporate email functions, or a combination of insourcing and outsourcing of a corporate security function all require different perspectives and tactical activities. Despite the commonalities among outsourcing functions (e.g., connectivity, user management, definition of task, measurements of effectiveness, goals and objectives), the same outsourcing blueprint for an inbound

call center would not serve for outsourcing corporate security. The objectives, rules, policies, risks, and rewards for each scenario are distinct and require customized attention.

Outsourcing decisions depend on which functions an organization decides it should perform for itself to maintain effectiveness and which functions would be performed more efficiently or more effectively by a vendor. Outsourcing also implies that an organization does not have, does not want to have, and cannot or will not have a specific expertise as part of its core business missions.

68.1.3 Insourcing. Insourcing is an accepted business practice; the government and many of today's large corporations hire contractors who work as insourcers within the physical and logical boundaries of their organization.

Insourcing poses risks to an organization because the place of work for the insourcer (vendor) is often within the physical boundaries of the organization—well within the physical perimeter, and in many cases, inside the logical perimeter defenses of the corporation. The contract worker is not an employee, but in most cases the contractor enjoys the same accesses as employees for the duration of the contract.

The defenses used for external protection are not as effective (if at all) against an insider with malicious intent. Security considerations regarding an insourced contractor require a different approach, more similar to layered security strategies employed for internal company resources.

Besides traditional IT security concerns, for these outsiders with insider access, the problems of human error, omissions, or complete bungling must be addressed. Security from external sources does not normally include internal human-error considerations. Before considering insourcing, an organization should already have a formal risk-management strategy (quality control) for internal human errors as well as for insider espionage and other insider threats.

An insourced contractor is not a full-fledged member of an organization and frequently resides somewhat outside of internal controls, thus requiring a separate and recognized classification and specific handling appropriate to the role. The General Accounting Office recently conducted a study of four federal agencies that rely on contractors to collect certain data on American citizens. The study found that "[a]gencies often do not limit the collection or use of information as required by the Privacy Act of 1974, … agencies don't ensure the accuracy of information … [and] contractors are not bound by those fair information practices and they often don't comply with all of them."[2]

Sound risk management requires acknowledgment and recognition that significant and substantial risks exist when insourcing. Attention to detail, a reality-based risk assessment, clearly articulated risks, attainable observations, and concrete audit points are all essential to manage successfully the insourced contractors and their mission. (See Chapters 13 and 45 in this *Handbook*.)

With the proliferation of mobile devices (smartphones, tablets) as well as independent data connections via USB data cards that connect to major telecommunications companies (e.g., ATT, T-Mobile, etc.), an insourced vendor is able to communicate externally without any oversight or examination of the content of those communications. Thus, trade secrets, competitive advantage information, future marketing plans, or even mundane operations details that could be examined and potentially exploited are free to leave the confines of the corporation without the knowledge of anyone in the corporation.

It might be prudent to expand our consideration of insourcing to include any situation where an individual who is not a direct employee of a corporation or entity is granted the ability to work within a corporate environment. This could include interns, volunteers, and even maintenance staff. With the economic downturn late in the first decade of the twenty-first century, many corporations have learned to rely on interns and volunteers who are hopeful that their efforts will be duly noted and potentially result in an offer of permanent employment. However, the flip side of such an arrangement is that the intern or volunteer could indeed be keenly observant and equipped with a smartphone, tablet, or laptop with a camera and an independent data connection; some of these temporary workers may make more money by selling a corporation's intellectual property, marketing information, or financial data to the highest bidder. (See Chapter 13 for more details of insider threats.)

68.1.4 Nearshoring. *Nearshoring* is the outsourcing of a specific discrete business function to a vendor located within the same, a nearby, or a bordering geographic region.

In some cases, an international outsource vendor will place components of its business in a specific country to acquire work within that country. For example, a United States–based outsource vendor will often position its operations in the United States to fulfill contracts from companies in the United States. Nearshoring can also include outsourcing to a bordering or regional country that shares a common cultural knowledge and understanding, as when a U.S. organization outsources to a group in Canada or Mexico.

Many of the largest technology companies in the world are some of the largest nearsourcing vendors. "Productized offerings from the large outsourcers include service desks, desktop management, and specialized network offerings. Examples include EDS Agile, HP SMB Services, and IBM Express Advantage."[3]

During the opposition to outsourcing that occurred in the first few years of the twenty-first century, the negative connotations of outsourcing were sometimes blunted when organizations nearshored their outsourcing work. The move toward globalization, and the painful lessons it taught certain nation-centric IT organizations, seemed easier to comprehend and accept if those jobs went to vendors located within relative proximity.

68.1.5 Offshoring. *Offshoring* is the outsourcing of specific, discrete business functions to a vendor whose corporate headquarters, or employees who fulfill the outsourced function, reside and work on another continent, as when a European company outsources to a company based in India.

Perhaps the most controversial kind of outsourcing, offshoring evokes a strong response, both in countries that move work offshore and in offshoring destinations. Nationalism, job security, self-interest, and a host of other emotions, both rational and irrational, seem to be part of every discussion. Those concepts, along with many others, are examined in Thomas Friedman's cornerstone book on outsourcing, *The World Is Flat*.[4] Friedman's perspective organizes the many disparate elements converged to "flatten the world" and to create a truly global work environment.

Regardless of the reasoning behind the decision, offshoring involves different considerations from insourcing and nearshoring. Fundamental business tasks taken for granted in a single geopolitical environment require different considerations when

offshoring. Many of the questions should be asked and answered long before addressing the ability of the offshore vendor and its employees simply to do the job.

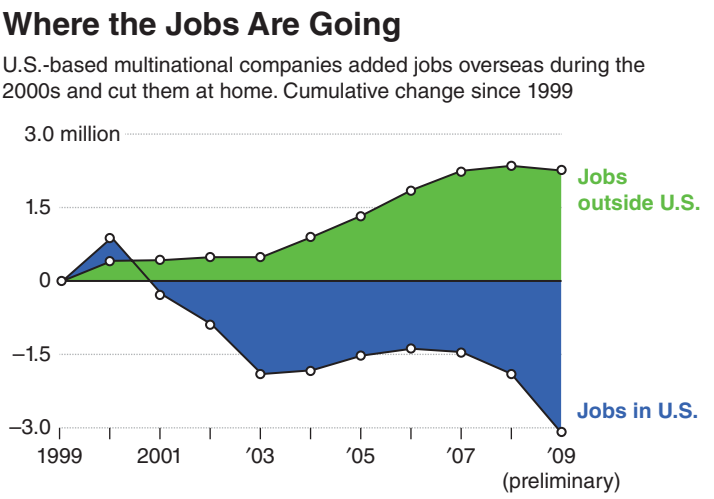
Despite all of these complications, offshoring is an established business practice and successfully accomplished every day. However, the rules that make offshoring successful will make any other business successful:

- Clearly understand your business.
- Clearly articulate your outsourcing (insourcing, nearshoring, and offshoring) goals.
- Possess a razor-sharp understanding of all risks, and managing those risks in an effective and reasonable manner.

68.2 WHY OUTSOURCE? There are two main drivers behind the growth of outsourcing today: the never-ending quests for greater organizational *effectiveness* and for greater *efficiency*. These drivers have come to the forefront of our economy because of the shift in strategic business thinking, begun in the 1990s, that is still affecting the way organizations are managing their businesses and serving their customers.

With the economic downturn that occurred late in the first decade of the twenty-first century, the impetus to outsource at a more rapid rate to low cost labor markets reached a fever pitch. In an effort to squeeze out costs within the corporation in many cases simply for economic survival, outsourcing became even more frenzied, bordering on the desperate.

In an article published April 19, 2011, entitled "Top 'US' Corporations Outsourced More Than 2.4 Million American Jobs Over The Last Decade," author Zaid Jilani states that contrary to trends in the 1990s where large U.S. corporations were creating jobs in the United States, these same large corporations "… have been adding more jobs abroad than at home," as is illustrated here:



Source: Zaid Jilani, "Top 'US' Corporations Outsourced More Than 2.4 Million America Jobs over the Last Decade," Think Progress Economy, April 19, 2011, http://thinkprogress.org/economy/2011/04/19/159555/us-corporations- outsourced-americans/?mobile=nc

Mr. Jilani's article continues:

> Another question asked of the executives found that the top reason for companies to outsource was to "reduce operating costs" (46 percent of respondents). Only 12 percent of respondents said their reason for outsourcing was "access to world class capabilities." This means companies are outsourcing to save themselves money, not to make better products.

In addition to sending jobs overseas, Mr. Jilani as well points out that many of the large corporations who have sent jobs overseas are as well attempting to repatriate the profits they've made overseas at a reduced tax liability. Mr. Jiliani gives the example of Cisco Systems.

> … which had 26 percent of its workforce abroad at the start of the decade but 46 percent of its workforce abroad by the end, is currently involved in a lobbying campaign titled "Win America" calling for a tax repatriation holiday that would let big corporations "bring money they have stashed overseas back to the US at a dramatically lower tax rate.[5]

68.2.1 Effectiveness versus Efficiency. Being *efficient* implies using an optimal amount of time and energy for getting a task done. In contrast, being *effective* means accomplishing the intended consequences of the task.

68.2.2 Being Effective. Historically, management has been primarily interested in increasing the efficiency of important business processes. But since the late 1980s, enlightened managers have realized that it can be very wasteful to try to optimize a process that does not lie within the core competency of the organization.

What are an organization's core competencies? A core competency is a *mission-critical* task or function that an organization is good at. "Mission-critical" refers to functions that are directly related to the strategic goals of the organization; for example, a hospital or a restaurant will consider cleanliness a mission-critical goal, whereas an automobile repair facility probably would not. Examples of core competencies may include reliable processes (e.g., Procter & Gamble's consumer brand management or Toyota's mantra of continuous improvement); a unique way of relating to customers (Nordstrom's superior retail service) or to suppliers (Boeing's supply chain management); or the particular look and feel of products or services (Apple's computers and iPods). A core competency also must meet these three conditions:

1. It can be used to develop entirely new products and services.
2. It provides significant customer benefits.
3. It is difficult for competitors to duplicate.[6]

If all three conditions are met, then a core competency can provide an organization with a true competitive advantage, which is a highly prized organizational asset, and which must be continuously guarded lest the competitive advantage be lost. A corollary is that all other functional areas, which are not within the core, are candidates for outsourcing.

By delegating tasks to vendors, an organization's management may concentrate more fully on its core business. Ideally, this will allow the organization to use its limited financial, talent, and other resources in the most productive manner. Knowing which activities an organization should perform, and which it should not, is the heart of being effective.

68.2.3 Being Efficient. Assuming one is doing the right things (i.e., being effective), the next logical question is whether one is doing them as efficiently as possible? Is the organization using the optimal amount of resources while still achieving quality and quantity goals? This question is important to both outsourcing organizations and to vendors performing outsourced work.

How can one measure efficiency? An often-sought-after outcome of efficiency is direct cost minimization. The idea of using outsourcing to achieve direct cost cutting is attractive. For example, according to a 2005 article appearing in *Mortgage Strategy*, British companies can save up to 40 percent by outsourcing all kinds of jobs to India. These companies report savings of a minimum £10 million pounds each year for every 1,000 jobs they outsource.[7]

In addition to direct cost reduction, other financial benefits of outsourcing to increase efficiencies include:

- **Decreased capital expenditure.** Because vendors utilize their own tools and infrastructure, there is no need for organizations to borrow money to purchase these items.
- **Decreased fixed costs.** Fewer fixed, lower, recurring payments allow an organization to manage its cash more easily because it would need less of it; fixed costs must be paid even when sales decrease, which may lead to a shortage of cash that could eventually bankrupt an organization.
- **Increased variable costs versus fixed costs.** If an organization pays a vendor based on units of work performed, then the costs of production become more variable, which is easier to manage as costs rise or fall based on the volume of goods or services required.

Another positive outcome of outsourcing is increased speed or work cycle time. By hiring a vendor located in India to test the quality of software written in the United States, an organization can perform software development work nearly 24 hours per day. Each morning on return to work, U.S. programmers would have the results of testing in India and could begin making corrections right away. By using this "follow the sun" approach, an organization can gain a speed advantage over its rivals.

Possibly the greatest aspect of efficiency that outsourcing can deliver is management focus. There can be great value in focusing a company's management team directly on those activities that differentiate it from the competition. Whenever a management team is focusing on noncore functions, it is usually operating not from a position of strength but from one of weakness. By definition, the team is not expert at noncore activities. In these cases, the team can spend too much time trying to understand and manage something that does not differentiate the organization from its competitors. If the noncore activities require too much management time and attention, there is a real risk that the core competencies of the organization may decrease in value. In the most serious of cases, an organization can lose its competitive edge completely, driving down sales, revenue, and profits.

68.3 CAN OUTSOURCING FAIL?

Yes, outsourcing can fail. For example:

> In 2004, J.P. Morgan Chase & Co. reassumed main technology functions following its merger with Bank One Corp., abandoning a US$5 billion pact with International Business Machines Corp. The same year, Electronic Data Systems Inc. abandoned a US$1 billion deal to run Dow Chemical Co.'s phone and computer networks.[8]

Deloitte Consulting's 2005 study, "Calling a Change in the Outsourcing Market," offers evidence that large organizations do not always achieve great efficiencies from outsourcing. The study, based on personal interviews with 25 of the largest organizations across eight industry sectors, reveals that:

- Seventy percent of participants have begun looking more cautiously at outsourcing after having negative experiences with outsourcing projects.
- One in four respondents have brought functions back in-house when outsourcing failed to deliver on promises of lower cost and more efficient operations.
- Forty-four percent of participants saw no cost savings as a result of outsourcing.[9]

It is important to keep in mind however, that for every outsourcing failure there are many other outsourcing successes. So then what separates successful from unsuccessful outsourcing/? Core to success are the following:

- A clear understanding of the corporation's business mission, and how the task to be outsourced contributes to that mission
- Rational and well-thought-out expectations of the outsourced function. If a corporation couldn't fully control the business function and continuously underperformed at the function, then don't expect outsourcing will be the panacea—unless there is corporate will (and dollars) to improve the function and the outsourcing includes hiring outside functional expertise for the function.
- Engagement with a world-class outsourcer. Unless a corporation has deep skills, knowledge, and insights in outsourcing, the risks of going with an untested or unknown outsourcing solution are simply too great. From infrastructure to data handling and security, there are simply too many moving parts for outsourcing to be a trivial task.

68.3.1 Why Does Outsourcing Fail? There are many risks to outsourcing, but direct cost reduction, often a top goal of outsourcing, can create the greatest risk of all. Indeed, "outsourcing deals most frequently stumble when they focus primarily on reducing costs."[10] Despite this observation, in a set of recent surveys, the rate of cost reduction as a driver for offshoring has been growing. From 2004 to 2006, the rate went from just over 70 percent of respondents to just over 80 percent.[11]

Total expenditure on outsourcing can meet or exceed the baseline of spending established prior to outsourcing. Although this may not be an inherently bad situation, management may perceive the experiment as a failure if the organization approaches outsourcing opportunities primarily to cut costs. However, if seen from a perspective of effectiveness rather than simply of efficiency, the outsourcing activities may be quite successful.

There are many other business reasons why outsourcing can fail. Looking more closely at the organization's perspective within the typical loss scenarios, in many cases the failure occurs because of one or more of these reasons:

- Monitoring and evaluating the performance of an outside vendor is a difficult task for many reasons, often due to inadequate selection and to monitoring of performance metrics.

- The outsourcing organization did not work hard enough to create a "win/win" opportunity for both itself and the vendor.
- Aligning the goals and priorities of the outside vendor with the organization's goals and priorities is rarely easy.
- Outsourcing requires an organization to master new and more complicated forms of communication, especially when outsourcing to a vendor in another country.
- It is challenging for an organization to ensure that the vendor's employees maintain sufficient knowledge and skill levels to do the work that has been delegated.
- Outsourcing usually introduces insecurity to the workforce and to the unions, as fear of job losses dominates their thinking. Terminating a contract usually requires a disruption in service continuity, and the quality of service can suffer, sometimes visibly to customers.[12]

Even when outsourcing appears to succeed from an operational perspective, there may be hidden inefficiencies in information protection that can decrease the overall value of the activity.

68.3.2 Intended and Unintended Consequences of Outsourcing. Now that the corporate world has experienced well over a decade of outsourcing, what in fact have we learned? In a lecture entitled "The Future of Outsourcing—Impact on Jobs," Patrick Dixon says:

> Outsourcing incentives are huge—and can lead to falls in service costs of 50–60%. Up to half of the $19 trillion spent every year by European companies on sales and administration could be outsourced.
>
> However, during 1995–2011 there has been a large migration of skilled, semi-skilled, and unskilled jobs from high-cost to low-cost nations. Outsourcing has produced labour shortages in many emerging economies. For example, China is now seeing 100% salary inflation at top end and India is not far behind—acute shortage of experienced business leadership. Some companies are now thinking of moving operations to places like Pakistan (50% lower costs and over 200,000 IT graduates looking for work), Bangladesh, or Vietnam. Changes are happening very quickly. But other corporations are already moving jobs back home—because cost savings have narrowed, and because outsourcing carries risks—in supply chain delays, language and culture challenges, distance management, loss of intellectual capital, and so on.[13]

Conversely, Mr. Dixon points out that many economists believe that while outsourcing does indeed move jobs to low-cost labor countries, there are subsequent higher-value jobs created in the original corporate country. Other cogent factors are:

- Outsourcing allows a company to control costs and by so doing stay in business.
- By staying in business the company can return value to shareholders.
- If the company remains economically viable it can as well live up to its obligations to its pensioners.
- By investing in low-cost labor markets, we are in effect helping those economies evolve and potentially become customers for our goods and services.[14]
- Reality, direct observation, and common sense tell us that moving from high-cost to low-cost labor markets creates the following cycle:
  - Technically skilled labor tasks move from high-cost (and experienced) labor markets to low- and mid-cost labor markets (with lower skills/less experience).

- The cost savings of moving from high-cost to low-cost labor can be significant on an hour to hour comparison, but because of the lack of experience and skill in the low-cost labor market, early savings are not as significant as expected. With less experienced labor, it simply takes longer to accomplish the same task as accomplished by experienced labor.
- Low-/mid-cost labor markets skills improve with experience. Those in high-cost labor markets that still remain focus on higher-level functions and much deeper technical expertise, and deeper skills transfer to the low-cost labor markets.
- As low-cost labor markets gain in experience and technical expertise, the high-cost labor teams grow smaller and smaller, and because the high-cost labor teams are no longer involved in the day-to-day technical operations, high-cost labor skills (especially experiential) begin an erosion cycle.

If we follow this cycle to its conclusion, given time, the technical and intellectual insights of a particular business now reside with an outsourcer, and the originating company has lost internal expertise to help lead and evolve the company into a viable future.

68.3.3 Universal Nature of Risk. Risk is inherent in virtually every human activity. One of the distinct advantages of an organized society is the ability of that society to distribute risk. Thus, not every member of society need manage every single risk.

When outsourcing, the greatest cost is that of ignorance, and the ultimate price is failure. Poorly defined expectations and poor planning resulting from a fundamental ignorance of a business will doom any corporate project, including outsourcing.

During early planning phases of the outsourcing project, it is prudent for the information assurance team to assess the entirety of the outsourcing project. However, information assurance does not have to own all aspects of the project; it is enough that the information assurance group should possess an end-to-end perspective, so that it can appropriately assess risk.

To relegate information assurance only to technical security tactics and practices is a serious error. Regardless of the vendor, the ultimate accountability (referred to as *responsibility*) stays with the firm. "One thing that can't be outsourced—responsibility .... Everything from employee policies to customer satisfaction to ethical and legal issues roots back to the impact on shareholder value. These responsibilities stay with the firm regardless of the functions that have been outsourced."[15]

The vendor, however, is accountable for carrying out all elements of the contractual agreement.

Thus, at the earliest planning stages, an information assurance review includes the overall scope of the outsourcing project. Questions should include:

- What information assets are at risk?
- What is the value and sensitivity of those assets?
- What current and future "risk shadow" will the outsourcing project cast on those corporate information assets?

The responses to these questions provide the foundation for the outsourcing project, as well as the degree of involvement from the information assurance community after that preliminary review.

Planning complex endeavors normally proceeds by careful identification and examination of concepts that move from general to specific. The sections that follow include broad categories for consideration at the onset of the outsourcing project. These categories require in-depth examination, as they relate to the specifics of each unique outsourcing project.

The outsource plan must include not only those tasks and security elements that are well known, but must as well consider the unthinkable. Amazon Corporation is an absolute pioneer and power in the cloud-computing world, providing hosting services to businesses large and small. As a cloud-computing provider, Amazon's Web services can have a direct impact on a business's revenue stream.

In an article entitled "Modern life halted as Netflix, Pinterest, Instagram go down" author Chris Matyszczyk describes how Netflix, Pinterest, and Instagram were completely unavailable: "The cause was reportedly an outage at Amazon's Elastic Compute Cloud in North Virginia, brought on by a thunder and lightning show."[16]

We need not look much further than very recent history to find other epic security incidents that involved corporations and outsourcers:

> …Sony shut down its PlayStation Network and mustc-streaming service Qriocity after an attacker made off with passwords and personal data from 77 million customer accounts and possibly 10 million encrypted credit card files. The information was in an AT&T data center in San Diego but controlled by Sony, which on Monday also pulled its Sony Online Entertainment site and Station.com gaming site after discovering an issue while investigating the intrusions.[17]

Alex Pham of the *Los Angeles Times* provides us a glimpse into the fiscal impact of the breach:

> In addition to losing an estimated revenue stream of about $10 million a week, Sony will probably have to reimburse customers who pay for its premium service, rebuild its computer systems, and beef up security measures, said Michael Pachter, an analyst with Wedbush Securities who said the incident could cost the company $50 million.
>
> That doesn't take into account the cost of defending against lawsuits, including two filed in California that are seeking federal class-action status.
>
> "Even if everyone loses just $1, that's $70 million," said Mark Rasch, director of cybersecurity and privacy at CSC, a computer networking and security firm in Falls Church, Va. Rasch, a former federal prosecutor who prosecuted Kevin Mitnick, a once notorious computer hacker, added, "If it's $2, you're looking at $140 million. Even lawyers can do that kind of math."[18]

Explaining those volumes of losses to a Board of Directors and shareholders is not an easy or pleasant task—thus a very clear understanding of risk and risk management in any outsourcing agreement is key. As the Sony incident as well as the Amazon–Netflix/Pinterest/Instagram examples demonstrate, even when dealing with world-class outsourcing leaders, bad things can and do happen; thus, a well-thought-out mitigation plan is not only essential, it is a matter of survival.

68.3.4 Clarity of Purpose and Intent. To outsource, one must possess the ability to articulate the task and to focus on creating a mutually beneficial vendor relationship. Poorly defined tasks lead to frustration and to unstructured attempts to meet inchoate needs instead of to measurable objectives—limitations that cause both organization and vendor to fail.

Corporations should associate with outsourcers that possess depth expertise in the particular business function intended to be outsourced. If a corporation outsources a

business function to an outsourcer that has to be told specifically how to fulfill the certain task, there is likely a disconnect. The purpose of outsourcing is to pay to have the function fulfilled by someone who can perform the function at more efficient and effective levels than the corporation itself.

Corporations depend on current employees to articulate the soon-to-be-outsourced task again, not the specific details of how they will attain the expected outcomes. Ironically, those employees best suited to articulate the task are frequently those whose jobs are most at risk after the successful implementation of the outsourcing project. This conflict of interest is a risk consideration.

Several phases of outsourcing include, but are not limited to, the collection and documentation of task knowledge for outsourced functions, vendor solicitation such as the request for information (RFI) or the request for proposal (RFP), vendor selection, and training.

Vendor selection is itself a key and integral part of the risk management of outsourcing, hence the importance of identifying, articulating, and quantifying the outsourcing goals. Once vendor selection occurs, corporate employees train vendor personnel on the outsourced functions. When the vendor personnel are trained and functional, most outsourcing results in the redeployment or release of corporate employees who fulfilled the specific task prior to outsourcing.

Risks related to clarity of purpose and intent include:

- Poor identification and definition of the outsourced task as well as clearly articulated success metrics results in an ineffective outsourcing program.
- Employees involved in the outsourcing project are aware that in many cases they are training their replacements.
- Any reduction in force regardless of the cause is a traumatic corporate event bringing unease and unrest to other employees.
- Unease and unrest can translate into lost productivity, poor morale, malaise, and corporate insider sabotage, all of which ultimately cause lost revenue.

Once outsource personnel are trained, especially if the technology/task is in high demand, the vendor must have safeguards in place to stem their attrition caused by their employees 'job hopping' to a more lucrative outsourcer.

68.3.5 Price. One of the largest risks of outsourcing is the formation of an unstable relationship with the vendor. When margins are paper thin, overall security practices often become the first victim. Recent findings from IT research firm Gartner indicate that a significant number of CIOs still look at outsourcing in terms of near-term profitability.

> As a result, they're setting themselves up for failure .... By 2008, more than 2.3 million offshore service workers will be employed by U.S. companies. But according to a recent Dun & Bradstreet survey, 20% of those outsourcing relationships will fail in the first two years, and 50% within five years.[19]

Risks related to price issues include:

- The vendor's culture must at least absorb, if not improve on, the originating company's corporate culture and do so in a fiscally responsible way.

- Defining outsourcing success by cost reduction alone introduces ongoing and increasing risk throughout the outsourcing relationship.
- Price and pricing are dependent on the ability to forecast business accurately. If the outsourced business does not have a perspective of historical trends, driving prices down at a vendor can actually have a negative effect on the outsourcer's outcomes and efforts.

Because outsourcers, like most businesses, seek to maximize profits in a low-margin industry, it is important to understand if the outsourcer actually outsources parts of their infrastructure (including security) to an even lower-cost outsourcer. For example, does an Indian outsourcing company outsource parts of their work (coding, reporting, training) to vendors in China, Vietnam, or other lower-cost labor markets? Understanding your outsourcer's business model and strategy is a key decision point.

68.3.6 Social Culture. The vendor must understand and address the social culture of the outsourced task and the differences from its own. Vendors must understand, absorb, and fulfill the outsourced task, under the same social norms as the sponsor organization. Cultures around the world have different societal norms, expectations, and nuances regarding confidentiality, possession, integrity, authenticity, availability, and utility. Vendors must be able to respond, react, and live within both the social culture they serve and their home culture.

Risks related to social culture include:

- Failure of the vendor to adapt to the social culture of the society served can result in lost money, wasted effort, potential legal issues, and customer dissatisfaction.
- Failure of organizations and vendors to anticipate, recognize, and mitigate problems arising from cultural differences regarding security, whether insourced or outsourced, poses a potentially catastrophic risk.
- Failure of the outsourcing corporation to understand the cultural nuances of the intended outsourcing country's social norms or the cultural nuances of the countries served by the outsourcer can drive out any potential gains of outsourcing with greater dissatisfaction in the outsourcing company, the outsourcer and the customer segments served

68.3.7 International Economics. The organization not only needs to understand its corporate and customer economics; it also must have a very deep and forward-looking view into the economic horizon of the vendor's geopolitical economy.

There are many willing vendors in emerging and low-cost labor markets. Without a clear understanding of the economic future of the vendor's country, a corporation can easily find itself tied to a vendor in an eroding or imploding economy.

Risks related to economic issues include:

- When the economy of a country collapses, any contractual relationship to an outsourcing vendor in that country will not survive.
- Forced into an evanescent relationship, a corporation may not have the ability to pull the outsourced task in-house, or to another vendor, quickly enough to save its own business.

In the second decade of the twenty-first century, China's economy contracts and wage growth increases, several European nations and the EU itself teeter on the verge of insolvency, and Vietnam and the Philippines see robust growth. Companies with long-term outsourcing contracts in highly changing regions may find themselves contractually locked into a difficult business climate for years to come.

68.3.8 Political Issues. Regardless of whether the vendor is nearshore or offshore, the client organization must have clear understanding, perspective, and acceptance of the political nuances in the vendor's country.

Political nuances and practices differ in every country. Acceptable behavior in one country may be reprehensible in another. Corporations must consider how political nuances intertwine with corporate objectives and policies.

Corporations must also be aware of and acknowledge the new risks and instabilities of the new century, and factor them into every nearshoring and offshoring outsourcing effort.

> The terrorist attacks Tuesday on trains in the western India city of Mumbai appeared unlikely to dampen investments and outsourcing to India … . The government also found evidence from terrorists killed in an encounter last year that they were targeting India's successful outsourcing industry.[20]

Risks related to political issues include:

- Are graft and bribery accepted or expected within the local political system? If so, how will such practices accord with national restrictions on the contracting organization? Will the outsourcing entity need to sponsor or participate in the political system of the vendor's country? How will such involvement be perceived nationally and internationally?
- Is the country politically stable? What are the opposing forces, and how do they see outsourcing vendors?
- Does the nearshoring or offshoring country possess an environment compatible with terrorist activities? What assurances of business continuity can the vendor make under such circumstances?

In mid-April 2012, the president of Argentina, Cristina Fernandez de Kirchner, announced "… a government plan to take YPF back through legislation that would expropriate a 51% share of the company. Through a decree, Fernandez immediately placed the leadership of YPF in the hands of Julio De Vido, the country's minister of planning and public investment."[21] Thus, the Spanish oil company Repsol lost its investment and purchase of Argentina's YPF, an oil and gas company. If a company outsources into a country where government seizures are acceptable and expected, the outsourcing company faces many risks related to infrastructure and, perhaps more importantly, intellectual property.

An energy crisis in India is having a significant impact on the Indian people and businesses. Contributing factors to the energy crisis include India's dependence on Iran for oil, yet U.S. and European sanctions are diminishing Iran's ability to deliver fuel to India. The diminished flow from Iran to India, as well as competition for oil and other fuels with China, Japan, and South Korea, leads commentators to suggest, "At stake is India's ability to bring electricity to 400 million rural residents—a third of

the population—as well as keep the lights on at corporate office towers and provide enough fuel for 1.5 million new vehicles added the roads each month."[22]

68.3.9 Environmental Factors. Site selection is one of the fundamental building blocks in any outsourcing project and a topic worthy of careful examination by itself. (For additional details, see Chapter 23 in this *Handbook*.)

Risks related to environmental factors include:

- If nearshore or offshore, will the vendor's regional infrastructure predictably and reliably support the expected and required service levels for the outsourced task?
- How susceptible is the nearshore/offshore location to natural disaster?
- Could a natural disaster trigger debilitating political or economic events in the host country?
- Does the host country possess the ability to recover from an environmental disaster?
- Could the outbreak of an easily transmitted disease or illness affect the location? And does the host country have the medical, technical, and communications infrastructure to deal with a pandemic?
- Is the area prepared for the actual weather it receives? For example, it is well known that the monsoon season comes to India. One of the causes of absenteeism in monsoon season in Mumbai is street flooding.
- Does the proposed outsourcing location have the external technology infrastructure that would support a remote workforce? If an employee cannot go to work because of street flooding or a pandemic, does the outsourcing country's telecommunications infrastructure support continuity of operations?

68.3.10 Travel. Whether nearshore or offshore, outsourcing will likely, and with regularity, send employees to the vendor for training, quality control, and other management functions. Important travel considerations focus on the costs of travel as well as on employee safety and health, addressing such issues as travel safety, food safety, medical preparations (immunizations, malaria pills, etc.), and locally available medical care.

Risks related to travel include:

- Will foreign workers be able to obtain required travel documents (e.g., exit permits from their home country, working visas for the outsourcing organization's country), allowing them to work in the client's organization?
- Extensive travel is grueling, takes a toll on employees' personal lives, and can cause productivity and management issues.
- Travel is expensive. Failure to budget for appropriate amounts of travel to the vendor will result in budget overruns.
- In the post-9/11 world, the risks and stresses to travelers have increased. Depending on the vendor's location, exposure of corporate employees to terrorist threats, or the fear of terrorist threats, is a significant consideration.
- With the advent of airport body scanners, controversy still occurs as to the long-term health risks of multiple scans by frequent travelers

- Keeping employees on the road traveling for long periods of time leads to physical, mental, and family stresses, all of which can contribute to employee turnover.

Common sense reveals that in most companies, it is the more senior and experienced personnel who do in fact travel in support of outsourcers. However, by traveling, those senior personnel face more and more nonproductive hours in airports and airplanes. Given that more and more major airlines have wireless availability while in flight, productivity can increase. Conversely, the ability to eavesdrop and intercept corporate communications increases in public places like airports, and in a small, enclosed, and public Wi-Fi network on an airplane, the chances for corporate espionage increase significantly.

68.3.11 Labor. Labor and workforce risks occur everywhere from Detroit, Michigan, to the Philippines and Bangalore, India. Without a clear understanding of the risks particular to the vendor's region, a corporation could find itself in a quagmire of constant turnover, escalating wages, work stoppages, and unfettered cost growth.

Careful examinations of current and forecasted workforce trends are core to the vendor selection process. As outsourcing sites gain in popularity, wages escalate driven by competition for workers. Once a corporation moves key functions to a vendor, if workforce conditions negatively change, previous cost savings could be lost. Important considerations include current and forecasted worker supply, the vendor's ability to retain employees (as evidenced by annual turnover rates), a propensity for collective bargaining in the vendor's country, and the stability of the government and the opposition's tactics.

Additionally, understanding the history of strikes or work stoppages in the vendor's location is a necessary preparation. For example:

> Political strife led to a shutdown Oct. 4 of most major outsourcing companies in Bangalore .... In April, Bangalore shut down for two days when citizens rioted following the death of Indian film icon Rajkumar. Published reports said the country's software companies lost $40 million in revenue.[23]

As mentioned in Section 68.3.7, the economic conditions in the offshore location should be followed carefully. Some countries suffering inflation may push workers toward collective action. For example, in April 2008, 20,000 Vietnamese workers at a factory making shoes for Nike went on strike to demand increased wages to keep pace with inflation.[24]

Another issue is whether worker exploitation (e.g., wages below a reasonable minimum, child labor, slave labor, unhealthful working conditions) exists in the offshore location or is practiced by the outsourcing vendor. Do such practices pose a threat to morale and reputation of the outsourcing organization?[25]

As we examine years of outsourcing experience, it is not uncommon for those workforces originally involved in lower-skill call-center positions to acquire experiences and skills over time which should lead to better paying and more challenging positions. However, if those more challenging positions don't materialize in an outsourcing location, the workforce can become disillusioned and migrate to another environment where the potential to advance both career and salary are possible.

68.3.12 Intellectual Property Risks. Outsourcing can and does include work which itself contains corporate intellectual property. By exposing intellectual property

to the outsourcer, that intellectual property can be easily copied and stolen. The U.S.-China Business Council noted,

> According to the release, six product areas are eligible for National Indigenous Innovation Product accreditation: computing and application hardware, telecom hardware, modern office equipment, software, new-energy products, and highly efficient energy-reducing products. These are the same six product areas outlined in the Notice Regarding the Launch of the National Indigenous Innovation Product Accreditation Work for 2009 (Circular 618), which was released in November 2009 … .
>
> Though the revised accreditation requirements are a step forward, several significant issues in China's indigenous innovation policies remain. In particular, the notice does not address the use of the product list or its link to government procurement preferences. Moreover, the proposed changes to the National Indigenous Innovation Product qualification criteria do not clear up questions about the relationship between the national product list and the continued validity and use of provincial- and local-level product lists that have been compiled based on discriminatory accreditation criteria. Finally, though the accreditation criteria no longer mention import substitution as a policy goal, the notice apparently does not change the November circular's application form, which asks whether an applicant's product can substitute for imports.[26]

In 2012, commentators reported that,

> It's no secret that China has been pressing multinationals to transfer more technology for accessing its huge market. Beijing now is turning another screw in its localization bid: requiring Chinese to run local operations of some foreign businesses.
>
> China in May ordered the Big Four auditors, Deloitte, PricewaterhouseCoopers, Ernst & Young, and KPMG, to have Chinese nationals manage their local units in three years … .
>
> As part of its "indigenous innovation" drive, Beijing in 2010 ordered that the technology of foreign firms must be registered in China, developed in China and, in some cases, made in China … .
>
> "We're fighting enough ghosts in China trying to protect our IP without having to intentionally share it," said Thomas Moga, an attorney at Shook, Hady & Bacon's Intellectual Property Section in Washington D.C … .
>
> While many emerging markets impose local content rules, none boasts China's economic weight. If Beijing continues to press, it could have a big impact on U.S. companies that rely on China for sales and manufacturing.[27]

Because foreign governments may have differing perspectives and practices around intellectual property, an assessment should be made of the financial consequences of losing control over intellectual property. A comparison of the size of the potential loss to the financial gains of outsourcing should be part of every outsourcing evaluation.

68.3.13 Additional Risks. Although the preceding lists may seem expansive, there are still many other fundamental and significant elements that must be part of any outsourcing decision-making process. Other complex and, in some cases, deeply fundamental considerations are:

- Loss of corporate expertise over tasks
- Loss of direct control
- Internal changes in corporate purpose (i.e., moving from a company of doers to a company that manages those who are doers)
- Overhead of ongoing contract management issues

- Currency fluctuations negatively impacting either the outsourcing corporation or the outsourcer
- Outsourcing staff turnover/retention
- Training (both failure to train, and failure to maintain training)
- Visa policies and the immigration of outsourced, highly skilled personnel to countries with permissive immigration policies for skilled labor
- Corporate espionage either directly or indirectly at an outsourcer

Understanding all of the risks allows the prepared organization to knowingly accept, mitigate, transfer, or ignore risks associated with the outsourcing project.

68.4 CONTROLLING THE RISKS. Outsourcing is an area in which the motto "Security transcends technology"[28] holds particularly true. Almost none of the threats to information itself, or information technology, is unique to an outsourced environment, although perhaps some may grow more dangerous in far-removed or foreign locales. Outsourcing does involve some serious security issues, but most—at first glance—might appear indirect threats to information assurance.

68.4.1 Controls on What? Most of the security issues of outsourcing involve people, corporations, societies, and governments. Security controls to mitigate the risks of outsourcing have little to do with technology such as computers and a great deal to do with organizational behavior. Since contractors perform crucial tasks but often are geographically far removed from those ultimately accountable for the tasks, the policies, contracts, agreements, and trust relationships that the organization has set up in advance will dictate the success of the endeavor. In particular, the organization's information-technology security policy takes on a much more visible role. The technological controls involve the assurance of interpersonal notions like trust and accountability, and are perhaps overshadowed by concern with legal matters, site selection, contractual obligations, politico-economics, and the impact of social networking and separation of duties.

68.4.2 Controlling Outsourcing Risk. Many of the risks described in Section 68.4 transcend the boundaries of individual risk types but are similar to other risk areas in the way they map to the six foundation elements of information security (the Parkerian Hexad).[29] It seems appropriate, then, to couch the discussion of controls in terms of the security foundations. This section focuses on controls that mitigate, for instance, confidentiality concerns in outsourced environments, touching on how the controls might affect the risk classes differently. Since part of the scope of this chapter is the outsourcing of security functions, some of the controls mentioned do not immediately appear to relate to information assurance but may affect other success metrics.

68.4.3 Availability Controls. When resources are local, the primary source of problems that can lead to a loss of availability has to do with physical and logical infrastructure. When resources are remote, the infrastructure issues are still important, and can indeed be more problematic. A good example of this difference is the February 2007 trans-Pacific cable cut that made much of Asia inaccessible over the Internet.[30] The farther the vendor from the home company, the more difficult and expensive it can be to acquire an alternate route for communications.

On June 29, 2012,

> …[L]ightning in Virginia took out part of Amazon's cloud computing service, called Amazon Web Services, which hundreds of companies use for data storage and computation. Well-known sites like Netflix, Pinterest, and Instagram were not accessible for hours. There was little information for customers about what had happened, or even whether user data was safe.
>
> The interruption underlined how businesses and consumers are increasingly exposed to unforeseen risks and wrenching disruptions as they increasingly embrace life in the cloud. It was also a big blow to what is probably the fastest-growing part of the media business, start-ups on the social Web that attract millions of users seemingly overnight … .
>
> The weekend's disruption happened after a lightning storm caused the power to fail at the Amazon Web Services center in Northern Virginia containing thousands of computer servers. For reasons Amazon was still unsure of on Sunday, the data center's backup generator also failed … .
>
> It was at least the second major failure for Amazon in that area. In April 2011, a problem in Amazon's networking at a nearby data center took down a number of applications and popular Websites, including Reddit and Quora, for more than a day … .[31]

Mitigation of availability risk revolves around planning. Sound backup strategies and business continuity plans should already be in place for the organization; the outsourcing project should also have these plans, both for the vendor site and for the staff at the home organization responsible for vendor communications. Given the economic, political, and environmental concerns that could lead to total (and possibly permanent) unavailability of the outsourced site, a backup vendor should be in place for disasters. These kinds of controls would be appropriate for natural disasters, labor strikes, terrorism, and a variety of other risks to availability. With the advent of cloud computing and, when ideally implemented, the geographical diversity of cloud computing providers, business continuity can benefit from the cloud computing platform, provided there are redundancies of path and infrastructure. (For details of backups, business continuity planning, and disaster recovery, see Chapters 57, 58, and 59 in this *Handbook*.)

If the vendor is in a foreign country where laws and contracts are enforced differently (or not at all), a service-level agreement incorporated into disaster planning might turn out to be a hollow, unenforceable contract.[32] Part of the evaluation of any outsourcing decision should be a visit to the site, including inspection of policy documents and physical tours of facilities. The organization needs to be sure that controls at the vendor site really do mirror the vendor's policy and contract documents and that the controls are instantiated contractually in an official statement of work. For example, it is a fairly common practice in outsourcing call centers that call center employees may not bring into the call center USB drives, recording devices, and the like. Yet, it has been these authors' first-hand experience that even in tightly controlled call centers, call-center employees routinely carry cell phones with gigabytes of memory and high-resolution cameras. Thus, not only is it important for controls to be contractually in place, but there must be the will and willingness on the part of the outsourcer to enforce the controls.

Beyond making sure that plans are in place against unavailability, the organization must be able to check up on the performance of the vendor. The agility required to switch to a backup vendor becomes much more possible with advance notice for at least some of the possible outages; however, this possibility must be evaluated with the knowledge that switching volumes or work between vendors is a complex task and should be field tested and validated on a regular basis. For instance, if the availability of a site relates to escalating economic problems causing workforce shortages, a periodic analysis of the regional news media, and of the vendor's work performance, might give

hints of problems on the horizon. Automated checking is appropriate as well, especially to keep track of network resources when access traverses the public Internet. Ultimately, for both people and technology, mitigating the risks of outsourced availability comes down to planning and monitoring.

68.4.4 Utility Controls. The utility of information (and of remote resources) hinges mostly on communication—both format and process. Careful version control avoids incompatible data. Encryption recovery agents can avoid the loss of utility if a user forgets a decryption password, but national and international restrictions on encryption must be kept in mind when planning the use of encryption across national borders.[33] These utility issues are common to any organization. With outsourcing, incompatible formats become more of a problem, especially with offshoring. If an application written in the United States uses ASCII encoding, but the vendor has applications that use Unicode, format issues can arise that need to be accounted for and solved. Although these problems may not become apparent immediately, planning for them must occur well in advance.

Human communication is also an issue with outsourcing. Even if the organization's native language is spoken by employees in the outsource vendor's site, it may not be their first language. Spoken communication (e.g., at a helpdesk) can suffer greatly if a technician's accent is too difficult for employees and customers of the client organization to understand. Similarly, written documentation and regular reports can lead to misunderstandings if language skills are not adequate. Either of these situations becomes manageable through advance planning for workforce training and through onsite liaison from the home organization, especially early in the relationship.

Utility of information also hinges on the type of data and the location of the data. Some government entities around the world require that data particular to that government or government constituents reside only within that country. Although the information may indeed be held within a specific country, there are additional considerations if those accessing the information are not within the country boundaries. Thus, the information may indeed be useful, but may not be useful to all based on their location.

68.4.5 Integrity and Authenticity Controls. The risk that the organization's data might be changed unknowingly, or replaced with other data, hinges on trust. Any time crucial business functions are given to an outside entity, trust issues arise. The organization must understand how much information is being shared, and with whom. Role-based access control and the principle of least privilege are appropriate here: Based on its role of supporting a particular business function, what is the minimum amount of privilege the vendor needs to do its job? Nevertheless, to be an effective support organization, the vendor may need access to corporate information that could prove damaging if misused.

From the perspective of integrity and authenticity controls, the decision to outsource must be accompanied by a decision about levels of trust. This trust should include an analysis of the vendor's history and reputation as well as a visual inspection of the site. The "trust infrastructure," which would include access-control mechanisms as well as division of labor and delegation of responsibilities, must be designed by the organization. Importantly, the vendor must not be able to make changes in this structure. Changes to trust relationships must be driven by the organization.

Given fluctuation of economies, job markets, and international relations, the decision to trust the outsourced personnel and processes should not be a one-time event. Ongoing

monitoring could reveal occasional lapses that might grow into bigger problems. How easy is it for the home organization to check up on the integrity and authenticity of data? Where are logs kept? Are backup copies of the logs (or the originals) sent to the home organization? What about change tracking on servers? How do technical personnel troubleshoot deeply technical issues that require access to protected information? Are all accesses of protected information logged and audited? What happens to protected customer information used in troubleshooting after it is no longer needed? Each of these questions should be addressed, and the answers should be written into the contractual language.

68.4.6 Confidentiality and Possession Controls. Merely making the decision to outsource partially compromises the confidentiality and possession of corporate information, just as the strength of a secret is decreased as soon as it is shared—even with a trusted confidante. The home organization must decide whether the loss of confidentiality and possession is balanced by the benefits of outsourcing. Within the United States, and in many countries with strong legal systems, laws protecting physical and intellectual property can help support this decision. The penalties enforced by the legal system serve as a deterrent to thieves and also serve to compensate damages in the event of a successful compromise. In countries where laws protecting intellectual property are weak (or absent), or where the rule of law is nominal or absent, this level of deterrent and compensation is not available, and the balance of risk shifts.

In June 2006, an employee at the HSBC bank call center in Bangalore, India, was arrested and charged with hacking into the bank's computers, breaching confidentiality agreements and privacy laws, and helping to steal £233,000. The accused was discovered to have been hired on the basis of forged school transcripts. According to news reports, the only criterion for hiring personnel into that call center was English-language skills.[34]

In the absence of strong legal backing, the organization can replace some of the deterrent normally provided by laws with language in the contract linking contractual compliance with payments and ongoing business relationships. The vendor should be required to meet security expectations, and the contract should specifically state that parts of the agreement (or the entire agreement) might be voided if security proves inadequate. Coupled with ongoing monitoring to catch problems before they become habitual or endemic, contractually tying security performance to the future of the relationship might prevent a damaging, large-scale loss of confidentiality or possession. And, although this kind of contractual language might not help in a foreign court, it should help protect the organization if the vendor manages to bring suit against the organization in the organization's home jurisdiction.

Outsourcing contracts should as well include specific requirements on the use of personal, portable electronic devices (smart phones, tablets, laptops, cameras) all of which are capable of independent data connections outside the purview of an outsourcer's technical infrastructure.

Additionally, rules regarding data exposure through social and professional networking sites should be written into outsourcing contracts. The personal use of sites such as Facebook or LinkedIn and the types of information that can be exposed by outsourced employees should be carefully considered as part of confidentiality and possession considerations in outsourcing contracts.

There are laws in the United States and other countries where governments may have access to confidential information when appropriate. Consider, for example,

cloud-computing outsourcers operating with a global footprint and managing data for a European company, in a European data center, yet managing the data center in the United States using U.S. employees. In those situations, the European companies may not understand U.S. laws; because they are used to the stricter privacy controls defined by the European Privacy Directive, they may think that the U.S. Government could have easy access to European corporate data because of U.S. laws. Similar issues arise for U.S. corporations with respect to other nations. The outsourcing provider and the outsourcing client must clearly delineate the impact of local laws the countries involved in the business relationships.

68.4.7 Making the Best of Outsourcing. Controlling the risks of outsourcing any function involves planning and careful implementation, primarily focused on trust and monitoring. The advice to *trust, but verify* applies particularly to outsourcing situations, in which a vendor necessarily gains at least some level of access to the organization's internal information and systems. Planning for risks to availability requires adequate business continuity and disaster recovery planning. Training, liaison, and careful planning well in advance of the outsourcing move are required to mitigate risks to utility. Integrity and authenticity controls, as well as confidentiality and possession controls, hinge on monitoring and enforcement, which can become problematic in different legal climates. Making contracts include business consequences for falling short of security requirements can help control the shortfalls of foreign jurisdictions. The ability to quickly recognize and contain evolving risk is key to long-term success in outsourcing.

68.5 OUTSOURCING SECURITY FUNCTIONS.

Delegating security functions to an outside vendor can increase the quality of an organization's overall security posture. This is done by leveraging the vendor's security expertise and perspective, which presumably it has acquired by providing a number of in-depth services to a large number of organizations.

Despite the media attention surrounding the outsourcing of some IT security functions, the use of a contracted guard force by organizations has been a common practice for years. As mentioned in Section 68.1.3, this is an example of insourcing: the use of contract or noncompany employees to fulfill certain business functions within physical and logical corporate boundaries.

In the last few years, organizations such as financial institutions have been nearshoring complex and costly IT security functions. This work is challenging for any organization, as the goal is to guard production networks against a never-ending stream of continuously changing threats. Also, these security functions usually do not pass the core competency tests described in Section 68.2.2, meaning they are good candidates for outsourcing. A leading example would be 24-hour-per-day monitoring and management of firewalls and intrusion-detection systems.

A new twist on a mature outsourcing tactic is offshoring software testing for security vulnerabilities. Assuming the vendor has the necessary tools and talent, outsourcing this function appears to make sense for many of the same reasons that organizations outsource quality assurance testing for any software development project. Not only is it an opportunity to gain effectiveness, but the efficiencies (e.g., faster cycle time and lower cost) can be compelling as well.

68.5.1 Who Outsources Security? A May 2013 report discussed the growth in security-outsourcing services:

> Despite some tumult, IT security outsourcing has been growing slowly but steadily during the last three years, with 38% of organizations outsourcing at least some element of this function in 2012. This growth will likely continue. More than one-fifth of survey respondents expect their IT security outsourcing to increase in the next year. At the same time, our outsourcing study shows strong cost benefits and service benefits to security outsourcing. In fact, the percentage of organizations that find outsourcing improves service is very high at 94%, tied for first place among the 11 IT functions included in our survey.
>
> Clearly, improving service is a strong factor motivating organizations to outsource their IT security functions. Why aren't more enterprise organizations doing so? Organizations are taking a measured approach to outsourcing the function because it is critical to the enterprise. The highly competitive outsourcing landscape could also be sowing uncertainty. Outsourcing partners need to do a better job of demonstrating their usefulness before the outsourcing of this vital IT function gains wider traction.[35]

68.5.2 Why Do Organizations Outsource Security? According to Levine, the top two reasons why organizations outsource network security are as a strategy for dealing with staffing challenges (effectiveness) and a desire for financial savings (efficiency).[36]

68.5.2.1 Staffing Challenges. Getting the most value from a security staff is an enduring challenge. Not only must talented people be found and retained, but they must be deployed so as to gain the greatest benefits. The highest value work for a security team includes those tasks that must be done from deep within the business context, such as policy setting, architecture, design, and risk management. These tasks require in-depth knowledge of a business's strategies, strengths, weaknesses, organizational structure, and culture. They are the core competencies of the corporate security team.

Given this situation, one of the best ongoing exercises a security officer or manager can perform to maximize the value of their team is to ask regularly, "Are my people exclusively working on tasks that cannot be delegated to anyone else?" Another way to ask this is, "Are my people doing any tasks that someone with less experience and organizational knowledge could perform at the same level of competency?"

Generally, such tasks are *not* mission-critical; they are *infrastructure* duties that support the key, distinguishing aspects of the organization's mission.

An example of a situation that can benefit from outsourcing security is protecting the network in an organization with limited technical resources.

Protecting the organizational network is a demanding job. Securing network connected digital assets (i.e., information systems) requires robust defenses against malicious hackers, viruses, worms, spyware, keystroke loggers, and denial of service attacks, just to name a few specific online threats. It may be difficult or impossible to find and employ enough security staff members who are both effective and efficient at dealing with all these threats. As well, threats and risk don't simplistically occur only during business hours. Maintaining adequate staff to manage and monitor $24 \times 7 \times 365$ is both physically and financially challenging.

In addition, understanding and defending against the latest threats requires constant education of staff, proactive monitoring, maintenance, and patching of the organization's network defenses. The capital expenditure for the care and feeding of specialized software and equipment can be very high.

Due to the nature and rapid pace of the threats, and a chronic shortage of talented people, there is an ongoing need to replace security staff. This adds additional expense for contract workers, decreased productivity, and increased risk of mistakes. Keeping information assurance specialists around after you have trained them is not easy; other organizations want and need them too. Outsourcing offers one way to transfer the burden of staff training and retention to another organization.

Finally, security service providers offer competent handling of routine security activities (i.e., monitoring and maintenance of hardware and software), and they can prepare the many reports required to document compliance with corporate policies and outside regulations. With these tasks being taken care of by a vendor, organizations can focus their internal efforts and personnel on more critical, high-value IT security functions.

68.5.2.2 Financial Savings. Despite our admonition in Section 68.4.4 to avoid outsourcing as a primary means to achieve direct cost reduction, a managed security service provider (MSSP) can offer tremendous economic efficiencies. For example, full-time security monitoring can be outsourced for significant cost savings.

A substantial challenge for all but the largest organizations is monitoring IT security 24 hours per day, seven days a week. For financial institutions in particular, this level of monitoring has become the de facto standard of due care. But providing that kind of constant vigilance is nearly impossible for many small organizations, such as community banks.

To provide 24-hour-per-day coverage with internal staff, organizations have to hire at least three full-time professionals, but would likely require twice as many to prevent staff burnout and a high turnover rate over the long term. Having backup coverage in place would require even more employees.

Managing an average week of 24-hour-per-day monitoring and response takes a minimum of five fully trained people. To begin with, you need one person for each of the three 8-hour shifts during the week. For weekends, the most economical approach is to have one person for each of two 12-hour shifts. Realistically, due to sickness, vacation, holidays, training, and other demands on staff, there would need to be between 8 and 10 employees to provide reliable coverage. In addition, should an actual security event occur, with the minimal coverage model as described here, there simply will not be enough hands to manage and contain the event. Thus, the affected company may indeed appreciate the scope and impact of the event, but be quite limited in their response.

The economics of this situation are straightforward. In the United States, a nearshore vendor can be hired to manage a small organization's firewalls, as well as run host intrusion detection and perimeter intrusion detection, for between US$25,000 and $50,000 a year.[37] In contrast, any organization in the United States would likely spend at least six times that amount on salaries alone to perform those functions with minimal staff. It is common for a highly skilled, in-house security professional to be paid between US$70,000 and US$180,000 a year in salary alone.[38] Add money for benefits (20 to 50 percent of base salary) as well as facilities, equipment, and other employment costs (80 to 150 percent of base salary), and the fully burdened cost for each employee can be almost two times what they are paid in direct salary.

At one of our organizations, the fully burdened cost of all employees is calculated every quarter by taking their direct salary and adding 184 percent.[39] Using this formula, the fully burdened cost of a security professional paid US$100,000 a year in direct salary is US$284,000.

68.5.2.3 Threat Intelligence and Additional Perspectives. An organization is best positioned to make decisions regarding IT security when it has the freshest, most complete intelligence (information concerning an enemy or possible enemy[40]) about emerging threats.

However, the tremendous volume of threat data being released every day can overwhelm all but the largest organizations. Even national governments struggle to keep up, and they usually have dozens or hundreds of people dedicated to gathering and analyzing data. Beyond problems of volume, there is the question of scope: It is difficult to know how any given vulnerability will be applicable to a specific system, if at all. A software maker's estimation of vulnerability severity is generic in nature; it may be more or less severe within the context of another organization. This is where an MSSP can really help.

MSSPs typically have the ability to gauge severity and spot trends based on what is happening to its other customers. One MSSP advertises that it

> … process[es] over a billion security events every day across more than 7,000 devices, giving our security research group unprecedented internal and external threat visibility across the globe. Using this visibility, [this vendor] maps the latest vulnerabilities and real-world threats to your infrastructure, enabling your team to prevent attacks.[41]

In Section 68.3.11, "the loss of corporate expertise over tasks" was mentioned as one of the risks of outsourcing. If your organization takes advantage of an MSSP's intelligence capability, and you have the opportunity and resources, a valuable provision in an outsourcing contract would be to transfer the vendor's intelligence-gathering skills to one or more employees of your organization. However, vendors may view such knowledge as proprietary.

68.5.3 What Are the Risks of Outsourcing Security? Trust is at the heart of the question of outsourcing risks. Can you trust the vendor to whom you are outsourcing, and its employees? Internal employees cannot typically gain the kind of intimacy with these people that they can with each other. This is a barrier to the human desire and tendency to build trust with others through direct interaction and observation.

If trust is at the heart of outsourcing security, the remainder of the vital organs consists of competency, capability, scope, and the rapid adaptability of the outsourcer. Trust is of course key, but the ability of the outsourcer to remain cutting-edge relevant in the light-speed world of security threats must be present.

There are specific risks in outsourcing security functions. *Total* control of an organization's security should never be transferred to an outside vendor. Although it may be possible to delegate (i.e., outsource) some operational duties, most companies find that keeping control of critical functions is vital to a successful security program. Examples of critical functions include firewall administration, direct control over all administrative/root accounts, and direct control over security logging. However, depending on your ability to manage the risk to your organization, you may feel comfortable delegating any of these functions. At the end of the day, outsourcing IT security is a very personal decision.

Some specific risks for IT security outsourcing include:

- An intrusion-detection system (IDS) monitoring vendor is more profitable as the number of events it responds to decreases. If the vendor is also responsible for

tuning your IDS, will your vendor tune the devices to produce the least number of events regardless of the actual threat environment?

- How can you ensure that in a significant customer-wide event, your organization will not be ignored or placed at a lower priority by your vendor?
- Is the vendor paid by "event" rather than by alert? If the answer is yes, what safeguards are in place to monitor potential abuse?
- Where are the outside vendor's personnel located: onshore, nearshore, offshore? Each would require specific contract language to accommodate language, cultural, and legal distinctions.
- Is the vendor vigilant about the background and expertise of its personnel?
- Who is monitoring the activities and behavior of the vendor's employees when they access your organization's data?
- Some countries are known as hotbeds of corporate espionage. Do you have strong control of your source code to ensure that the vendor's employees do not share the code with any other persons or entity?
- What are the vendor's source country's governmental policies on privacy and intellectual property?
- What is the vendor's employee turnover rate?
- What training takes place for each of the vendor's employees? Is the training frequently updated and refreshed? Do existing employees have access to the training as needed?
- An IT security outsourcer obviously serves many customers. In the case of a large-scale event, where will your corporation rank in order of importance, urgency, and response? Will your IT security outsourcer first attend to their largest clients, and eventually get to your corporation?

68.5.4 How to Outsource Security Functions. Although not very different from outsourcing any business process, there are some unique aspects to outsourcing IT security functions. This section describes what is different and provides specific examples. We encourage you to consult with your contracting office or other reputable and knowledgeable sources for a more thorough treatment.

68.5.4.1 Where to Begin. As with any outsourced work, first gather your business requirements and use them to define the outcomes you expect the vendor to deliver. Normally, this information is delivered to the vendor as a statement of work (SOW). Include in your SOW only those functions that do not require intimate knowledge and experience of the specific, mission-critical functions of the organization. As discussed in Section 68.2, outsourcing should allow management to focus on its mission-critical functions; thus it is inappropriate to outsource functions that require great experience and insight in those areas.

The following discussion includes examples from a specific case study on outsourcing user-account management.

68.5.4.2 Brief Case Study, Part I. One author of this chapter recently delegated responsibility for resetting passwords and creating user accounts to an outside vendor. The business driver to delegate was the result of an analysis of effectiveness of his internal team: There was just too much work to do for the number of

employees authorized. In addition to asking whether a less experienced (i.e., less expensive) team could accomplish certain tasks, the security team also considered the risks to the organization if the work were not done correctly.

In a team decision meeting, we determined that resetting passwords and creating accounts for nonadministrative users was of sufficiently low risk that we could delegate those tasks. In contrast, we determined that performing these tasks on administrative and service accounts was too risky to delegate, so we continued to do that work. The fact that the daily volume of administrative account work was relatively low compared to nonadministrative accounts helped justify keeping those tasks in house.

68.5.4.3 Determining Desired Outcomes. When defining outcomes for operational IT security tasks, focus on the urgency and severity levels of the responses. This is typically driven by the response times required by the organization.

68.5.4.4 Brief Case Study, Part II. Once we decided to delegate resetting passwords and creating user accounts, we then asked operational business leaders what responsiveness they required. Based on their input, we assigned password resets as severity level 1, requiring no more than four hours to complete on receipt of request. User account creation is severity level 2, requiring no more than one business day to complete on receipt of request.

68.5.4.5 Choose a Reliable Vendor. Once your SOW is complete, choose a reliable IT security vendor. Consider these points before you make your final decision[42]:

- Financial health (e.g., annual revenues, longevity)
- Reliable infrastructure (e.g., state-of-the-art tools; disaster resistant)
- Competent staff (e.g., sufficient experience; appropriate credentials; access to global intelligence; original training plans; ongoing refresher training)
- Satisfied customers (e.g., talk with MSSP references; ask what they like the least about the MSSP)
- Vendor independence (e.g., no hidden financial motives to sell you solutions that are not appropriate for your organization; third-party attestations of effectiveness)
- Appropriate service-level agreement (e.g., contains requirements, outcomes desired, response times, roles and responsibilities, and metrics)
- Legal safeguards (e.g., recourse for your organization in the event of vendor breach of contract; if offshored, the applicable laws in the host country that might prevail in a contractual issue, or in a government "nationalization" situation)

68.5.4.6 Service-Level Agreements. Once a reliable vendor is chosen, a strong, precise service-level agreement (SLA) is required. The SLA must specify outcomes desired, response times, roles and responsibilities, metrics, and other requirements. Ideally, the SLA should be written to remain relatively stable over time. Realistically, an SLA needs reliable provisions for change and conflict resolution because the requirements will change over time, no matter how much homework you do.

SLAs need to be realistic given both the scope of the SLA and the financial underpinnings of the contract and statement of work. Demanding an SLA that could never be met, given contractually agreed-to staffing levels, would be useless.

68.5.4.7 Metrics. Include metrics in the SLA that will allow you to measure outcomes; install an independent monitoring system to validate reported metrics and to detect unauthorized behavior by vendor staff. Periodically (weekly at first, then monthly if performance warrants), review the metrics and activities with the vendor's management to ensure that work performed matches the requirements of the SLA; adjust the SLA as necessary to drop metrics that are not useful and add new metrics that are.

It is key to select performance metrics with great care. Reviewing metrics snapshots on a daily basis is important, but even more important is the trends in metrics over time. Many security incidents are not a "big bang," but instead are slowly expanding and growing "under the radar." Having carefully considered metrics that allow a window into both the vendor's adherence to the contract and a perspective of the health of the corporation are the ultimate goals.

68.5.4.8 Brief Case Study, Part III. We monitor the vendor response times for resetting passwords and creating user accounts by examining the time stamps in the vendor's ticketing system correlated with the events written in the system logs. To ensure the vendor does not tamper with the system logs, we have many events immediately forwarded to a centralized logging server to which the vendor has no access.

68.5.4.9 Gaining Maximum Efficiency. Unless you have a good reason not to do so, allow the vendor to determine how it will deliver your desired outcomes. This gives the vendor the ability to determine how to achieve maximize efficiency. Of course, a vendor is inherently motivated to do this in order to be competitive and maximize profits. Be aware that this also means the vendor has incentives to cut corners, which could result in security incidents for you.

Be clear on expected results while still allowing the vendor the flexibility to determine how they plan to deliver those results. However, you as the outsourcing corporation do have a responsibility as well to review and accept the final "how" the vendor will deliver to your result specifications.

68.5.4.10 How the Vendor Does the Work. Once the vendor has created its procedures, be sure to review and approve how the vendor will do the work. After all, security is based not just on what things you do but how you do them. The ultimate authority in the matter will be your organization's information-security policy. Be careful, though, not to criticize vendors just because they do not do things the way your team would. Your priorities are to get the right outcomes with the least risk at an affordable price. Be very careful not to upset your priorities over nothing more than personal preferences.

68.5.4.11 Brief Case Study, Part IV. In outsourcing password resets and account creation, we allowed the vendor to devise its own procedures. In our case, it made sense to permit this level of freedom because the vendor had tools and know-how that our team did not. For instance, we did not have an incident ticketing system, and the vendor already had one up and running. This allowed the vendor to receive, log, and assign work more efficiently than we could.

However, after some time, one of the risks we noted in Section 68.6.3.10 reared its ugly head. We found that the vendor sometimes placed passwords in the ticketing

system because doing so saved time (it did not have to reset the passwords every time it went into an account) and allowed the ticket to be closed more quickly. Clearly, we were not as motivated by their profit motive as by preserving our security posture. This is why we regularly monitored the vendor's ticketing system by having a console installed directly in our work area.

68.5.4.12 Holding the Vendor Accountable. Finally, be sure to include a strong penalty provision in your SLA to deter careless mistakes by the vendor. If your monitoring tools are effective, you will know about mistakes before the vendor does. Be thoughtful, though, because you do not want to deter the vendor from reporting those mistakes. The point of discussing a mistake should not be to punish as a first recourse, but to correct the vendor in the hope that it will grow in competence and become a better steward of your security work.

68.5.5 Controlling the Risk of Security Outsourcing. Organizations should consider their decision to outsource security in terms of their organization's overall outsourcing strategy, and should determine if their organization has the skill sets and tools necessary to manage the outsourcing relationship. Also, they must realize that contract management skills will be the primary management tool, not IT management expertise.

Ongoing monitoring by the outsourcer will be required. Organizations cannot take for granted the presence and effectiveness of monitoring activity. Provisions must be made to ensure that organizations get the services they are paying for. Should the vendor fail to deliver the services, the organization should be prepared to perform the work once again, quickly and effectively. Once confidence in an IT security vendor has been lost, the vendor's administrative access should be revoked as soon as possible.

Refer to Section 68.4 of this chapter, controlling the risks of IT outsourcing, for more advice.

68.6 CONCLUDING REMARKS. *Outsourcing* is a new term for an old concept: We humans have always liked to do what we are good at doing, relying on the skills of others to fill in the gaps in our competencies. We outsource tasks to concentrate on being more effective at our core skills, and we hope to become more efficient as a result. Types of outsourcing—insourcing, nearshoring, offshoring, cloud computing—differ based on how far from the parent organization the outsourcing vendor operates, but there are other much more significant differences when evaluating the risks of outsourcing.

Many organizations have saved significant amounts of money by outsourcing some functions both task and infrastructure. Hindsight, however, shows that many organizations that outsource solely to save money find themselves with significant problems caused by a lack of careful evaluation of outcomes and risks. Proper planning should include:

- Being clear about the purpose of the outsourcing
- Avoiding the tendency to focus only on cost
- Adequately understanding social culture
- Accounting for differences in politics and economics
- Analyzing potential problems in infrastructure and environment
- Anticipating travel and labor issues

Any organization that looks at all these issues well in advance of an outsourcing decision stands a good chance to succeed, despite the potential pitfalls.

The keys to mitigating risks in any outsourcing project—careful planning and constant vigilance—are particularly applicable if security functions are the subject of an outsourcing decision. Depending on the security stance and regulatory environment of the home organization, some security functions may be able to be outsourced in a way that does not put the organization at unacceptable risk. In the end, the great care required to outsource security functions properly has the potential to improve security throughout the organization, and perhaps even make security easier to do well at a vendor site than at the home organization.

68.7 FURTHER READING

Axelrod, C. W. *Outsourcing Information Security*. Artech House, 2004.

Carmel, E., and P. Tjia. *Offshoring Information Technology: Sourcing and Outsourcing to a Global Workforce*. Cambridge University Press, 2005.

Cohen, L., and A. Young. *Multisourcing: Moving Beyond Outsourcing to Achieve Growth and Agility*. Harvard Business School Press, 2005.

Cooney, M. "Can You Trust China for Outsourcing?" *Network World*, May 29, 2006. www.networkworld.com/columnists/2006/052906edit.html

Cullen, S. *Intelligent IT Outsourcing: Eight Building Blocks to Success*. Butterworth-Heineman, 2003.

Dara, N. "Cyber Crime Comes of Age as Foreign Plugs Sell Secrets," *DNA India*, October 19, 2005; www.dnaindia.com/india/6411/report-cyber-crime-comes-of-age-as-foreign-plugs-sell-secrets

Grance, T., J. Hash, M. Steven, K. O'Neal, and N. Bartol. *Guide to Information Technology Security Services.* NIST Special Publication SP 800-35 (October 2003). http://csrc.nist.gov/publications/nistpubs/800-35/NIST-SP800-35.pdf

Greaver, M. F. *Strategic Outsourcing: A Structured Approach to Outsourcing Decisions and Initiatives*. AMACOM, 1999.

Jansen, W., and T. Grance. *Guidelines on Security and Privacy in Public Cloud Computing*. NIST Special Publication SP 800-144 (December 2011). http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf

Kendrick, R. *Outsourcing IT: A Governance Guide*. IT Governance Publishing, 2009.

Koulopoulos, T. M., and T. Roloff. *Smartsourcing: Driving Innovation and Growth Through Outsourcing*. Platinum Press, 2006.

Power, M. J., K. Desouza, and C. Bonifazi. *The Outsourcing Handbook: How to Implement a Successful Outsourcing Process*. Kogan Page, 2006.

Sood, R. IT, *Software and Services: Outsourcing and Offshoring*. AiAiYo Books, 2005.

Stees, J. *Outsourcing Security: A Guide for Contracting Services*. Butterworth-Heineman, 1998.

Vashistha, A. *The Offshore Nation: Strategies for Success in Global Outsourcing and Offshoring*. McGraw-Hill, 2006.

68.8 NOTES

1. M. Cooney, "Outsourcing Bonanza 2006: 8 Outsourcing Trends You Need to Know About," *Network World*, December 13, 2006; www.networkworld.com/news/2006/121306-outsourcing-trends.html

2. M. Cooney, "Government Agency Outsourcing Firms Don't Respect Private Data, GAO Reports," *Network World*, April 5, 2006; www.networkworld.com/news/2006/040506-gao-outsourcing.html
3. Dan Twing, "10 Reasons Why Small Businesses Should Consider Outsourcing," *Network World*, July 5, 2006; www.networkworld.com/newsletters/asp/2006/0703out1.html?page=1
4. T. L. Friedman, *The World Is Flat: A Brief History of the Twenty-First Century* (Farrar, Straus, and Giroux, 2006), 48–172.
5. Zaid Jilani, "Top 'US' Corporations Outsourced More Than 2.4 Million America Jobs Over The Last Decade," *Think Progress Economy*, April 19, 2011, http://thinkprogress.org/economy/2011/04/19/159555/us-corporations-outsourced-americans/?mobile=nc
6. G. Hamel and C. K. Prahalad, "The Core Competence of the Corporation," *Harvard Business Review* 68, No. 3 (May–June 1990): 79–93.
7. "Tempted by the Call of the East," *Mortgage Strategy* (August 22, 2005), 3.
8. S. Thurm, "Behind Outsourcing: Promise and Pitfalls," *Wall Street Journal* (February 27, 2007).
9. D. E. Levine, "Farming Out Network Security: Outsourcing Might Save You Money If You Choose Your Provider with Care," *Security Technology & Design*, May 2005, www.highbeam.com/doc/1G1-133132842.html
10. D. E. Levine, "Farming Out Network Security."
11. Duke University/Archstone Consulting LLC Offshoring Research Network 2004 and 2005 surveys; Duke University/Booz Allen Hamilton Offshoring Research Network 2006 survey. As reported by Mary Brandel, "Offshoring Grows Up," *Computerworld*, March 12, 2007, www.computerworld.com.au/article/174894/offshoring_grows_up
12. D. Winkelman, D. Dole, L. Pinkard, J. Molloy, D. Willey, and M. Davids, "The Outsourcing Source Book," *Journal of Business Strategy* 14, No. 3 (May–June 1993): 52.
13. P. Dixon, "The Future of Outsourcing—Impact on Jobs," Global Change.com; www.globalchange.com/outsourcing.htm
14. P. Dixon, "The Future of Outsourcing."
15. D. Twing, "Reviewing the Security Aspect of Outsourcing," *Network World Outsourcing Newsletter* (September 7, 2005).
16. Chris Matyszczyk, "Modern Life Halted as Netflix, Pinterest, Instagram Go Down," C|Net, June 30, 2012, http://news.cnet.com/8301-17852_3-57464342-71/modern-life-halted-as-netflix-pinterest-instagram-go-down
17. Donna Howell, "Security, Outages Gray Areas for Cloud Computing," *Investor's Business Daily* (May 3, 2011), A5.
18. Alex Pham, "PlayStation Network Security Break Will Cost Sony Much More than Money," *Los Angeles Times*, April 28, 2011, http://articles.latimes.com/2011/apr/28/business/la-fi-0428-ct-sony-hack-20110428
19. K. Evans-Correia, "Outsourcing on Verge of Cultural Evolution," SearchCIO.com, April 18, 2006, searchcio.techtarget.com/news/1179791/Outsourcing-on-verge-of-cultural-evolution
20. "Mumbai Blasts Should Not Affect Investments to India," IDG News Service, July 12, 2006, www.itworld.com/Man/2701/071206mumbai/pfindex.html

21. CNN Wire Staff, "Argentina, Spain at Odds over Oil Company Expropriation," CNN, April 17, 2012, http://articles.cnn.com/2012-04-17/americas/world_americas_argentina-ypf-repsol_1_ypf-julio-de-vido-argentine-government?_s=PM:AMERICAS
22. A. Sharma and M. Bahree, "Grinding Energy Shortage Takes Toll on India's Growth," *Wall Street Journal* (July 2, 2012), 1.
23. M. Cooney, "No Outsourcing Today: Strike Closes India's Tech Center," *Network World*, October 5, 2006; www.networkworld.com/news/2006/100506-strike-closes-india-tech-center.html
24. "20,000 Vietnamese Workers Go on Strike at Nike Contract Factory," Associated Press, April 1, 2008, usatoday30.usatoday.com/money/industries/retail/2008-04-01-vietname-nike-strike_N.htm
25. A. Fallone, "Overview Child Laborers," IHS Child Slave Labor News, 2005, http://ihscslnews.org/view_article.php?id=54
26. The US–China Business Council, "China Proposes Partial Solution to Indigenous Innovation Issues; USCBC Seeks Member Comments," USChina.org Website, April 12, 2010, https://www.uschina.org/public/documents/2010/04/indigenous-innovation-memo.html (URL inactive)
27. Doug Tsurruoka, "China Goes 'Local' On US Firms' Units," *Investor's Business Daily* (July 2, 2012), 1, http://finance.yahoo.com/news/china-goes-local-u-firms-222400546.html
28. $(ISC)^2$, International Information Systems Security Certification Consortium: www.isc2.org
29. We owe a lot to Donn Parker for this new way to look at information assurance. See Chapter 3 in this *Handbook* for his description of these security elements.
30. H. Timmons, "Cut Cable Disrupts Web and Phones in India and Middle East," *International Herald Tribune*, January 31, 2008, www.nytimes.com/2008/01/31/technology/31iht-net.4.9648965.html
31. Q. Hardy, "A Storm's Disruption Is Felt in the Technology Cloud," *The New York Times*, July 1, 2012, www.nytimes.com/2012/07/02/technology/amazons-cloud-service-is-disrupted-by-a-summer-storm.html
32. See, for example, A. Eunjung Cha, "New Law Gives Chinese Workers Power, Gives Businesses Nightmares," *Washington Post*, April 14, 2008; www.washingtonpost.com/wp-dyn/content/article/2008/04/13/AR2008041302214.html
33. See, for example, J. Markoff, "Encryption Tool Rekindles Security Debate," *International Herald Tribune*, May 21, 2006, www.nytimes.com/2006/05/21/technology/21iht-privacy.html
34. K. V. Subramanya, "HSBC Scam: U.K. recipient of data is of Indian origin," *The Hindu*, June 30, 2006, www.hindu.com/2006/06/30/stories/2006063002920700.htm
35. "IT Security Outsourcing Shows Steady Growth," *Computer Economics* (2013-05), www.computereconomics.com/article.cfm?id=1831
36. Levine, "Farming Out Network Security."
37. C. Costanzo, "Internet Security: Outsource or Go It Alone?" *Community Banker* 14, no. 6 (June 2005), http://findarticles.com/p/articles/mi_qa5344/is_200506/ai_n21373311/pg_1

38. "Information Security Officer Salaries," Glassdoor.com, July 2013, www .glassdoor.com/Salaries/information-security-officer-salary-SRCH_KO0,28.htm
39. This calculation was current in March 2007. The formula described and the 184 percent figure are the result of our consultation with the finance department at one of the authors' employers. The name of the employer is being withheld to protect their confidentiality.
40. *Merriam–Webster's Collegiate Dictionary*, 11th ed., 2004, p. 650. Also online at www.merriam-webster.com/dictionary/intelligence
41. Dell SecureWorks, "Threat Intelligence Service," www.secureworks.com/services/threat_intelligence.html
42. J. Mears, "Is Security Ripe for Outsourcing?" *Network World*, August 23, 2004, www.networkworld.com/news/2004/082304outsecure.html

INTRODUCTION TO PART VIII

PUBLIC POLICY AND OTHER CONSIDERATIONS

This edition of the *Handbook* ends with compelling issues in information security. Part VIII provides a basis for vigorous discussion about important and controversial topics such as:

69. **Privacy in Cyberspace: U.S. and European Perspectives.** With increasingly frequent losses of control over personally identifiable information, the public is ever more concerned about privacy, especially with recent revelations of widespread NSA surveillance and analysis of electronic communications
70. **Anonymity and Identity in Cyberspace.** How individuals are representing themselves in Internet-mediated communications; the social and legal consequences of completely anonymous interactions, and of untraceable but stable identifiers
71. **Healthcare Security and Privacy.** How the special requirements of high availability coupled with extreme sensitivity of medical information poses complex problems for security specialists in medical environments
72. **Legal and Policy Issues of Censorship and Content Filtering.** How corporations and governments around the world regulate access to information that violates social norms or is perceived as a potential threat to state power
73. **Expert Witnesses and the *Daubert* Challenge.** How security specialists should prepare for their day in court
74. **Professional Certification and Training in Information Assurance.** Benefits and costs of education, professional certifications, examinations, and commercial training
75. **The Future of Information Assurance.** A brilliant young computer scientist offers his take on possible directions of interest in our field

CHAPTER 69

PRIVACY IN CYBERSPACE: U.S. AND EUROPEAN PERSPECTIVES

Henry L. Judy, Scott L. David, Benjamin S. Hayes, Jeffrey B. Ritter, Marc Rotenberg, and M. E. Kabay

69.1 INTRODUCTION: WORLDWIDE TRENDS 69·1
- 69.1.1 Current Cyberprivacy Issues 69·2
- 69.1.2 Laws, Regulations, and Agreements 69·5
- 69.1.3 Sources of Privacy Law 69·7

69.2 EUROPEAN APPROACHES TO PRIVACY 69·8
- 69.2.1 History and Organization for Economic Cooperation and Development Principles 69·8
- 69.2.2 European Union Data Protection Directive 95/46/EC 69·8
- 69.2.3 Harmonization of Non–EU Countries to the EU Directive 69·11
- 69.2.4 European Union Telecommunications Directive 69·11
- 69.2.5 Establishment of the European Data Protection Supervisor 69·11

69.3 UNITED STATES 69·12
- 69.3.1 History, Common Law Torts 69·12
- 69.3.2 Public Sector 69·12
- 69.3.3 Private Sector 69·16
- 69.3.4 State Legislation 69·27

69.4 COMPLIANCE MODELS 69·28
- 69.4.1 U.S. Legislation 69·28
- 69.4.2 U.S. Federal Trade Commission Section 5 Authority 69·29
- 69.4.3 Self-Regulatory Regimes and Codes of Conduct 69·29
- 69.4.4 Contract Infrastructure 69·29
- 69.4.5 Synthesis of Contracts, Technology, and Law 69·30
- 69.4.6 Getting Started: A Practical Checklist 69·31

69.5 FURTHER READING 69·31

69.6 NOTES 69·32

69.1 INTRODUCTION: WORLDWIDE TRENDS. As information technology continues to develop, increased opportunities and incentives are created to collect information about individuals and to use that personal information for diverse and lucrative purposes. Some believe that, in a global information economy, the most economically valuable electronic asset will be aggregations of information on individuals. For over

two decades, concerns regarding the privacy of the individual and cross-border movement of personal information have been reflected in international policy frameworks, such as the *Organization for Economic Cooperation and Development* (OECD) *Privacy Guidelines*. In recent years, these concerns have escalated in Europe, the United States, and other economies, particularly due to the rapid commercialization of the Internet and the development of new and more powerful information technologies. There are also significant developments with DNA databases, RFID technologies, geopositioning data, and electronic health records. These concerns have accelerated the growth of privacy law, with recent enactments in the United States, Europe, and other areas of the world.[1] At the same time, new policy frameworks have emerged, most notably in the Asia-Pacific region.

As this chapter was going to press in July 2013, cyberprivacy issues in the United States and Europe were actively being covered by the news media and generating concern among civil liberties groups. The domestic spying programs of the U.S. National Security Agency and questions about illegal phone tapping in the United Kingdom were top news items.

69.1.1 Current Cyberprivacy Issues.

Several cases have caused interest among privacy advocates in the first decades of the twenty-first century.

69.1.1.1 NSA Domestic Spying.

The Electronic Frontier Foundation (EFF) page about extra-legal surveillance of U.S. residents includes this succinct summary:

> In October 2001, President Bush issued a secret presidential order authorizing the NSA to conduct a range of surveillance activities inside of the United States without statutory authorization or court approval, including electronic surveillance of Americans' telephone and Internet communications. This program of surveillance continues through today, although the legal justifications have changed over time, and works with the major telecommunications and Internet companies.
>
> In 2005, after the *New York Times* broke the story of the surveillance program, the President publicly admitted one portion of it—warrantless surveillance of Americans believed to be communicating with people connected with terrorism suspects—Senior Bush Administration officials later confirmed that the President's authorization went beyond the surveillance of terrorists and conceded that the program did not comply with the Foreign Intelligence Surveillance Act (FISA). The President, invoking a theory of limitless executive power to disregard the mandates of Congress, reauthorized this warrantless surveillance more than thirty times, including after the Department of Justice found the program to violate criminal laws. President Obama has continued the program, but with differing secret legal justifications. Obama has given no public legal justification for it and, in some situations, appears to be strategically denying certain portions of it. For other portions, including the collection of telecommunications records, the Obama Administration said it could neither confirm nor deny its actions until May 2013, when the DNI finally admitted additional portions of it. Members of Congress have confirmed that additional domestic surveillance by the NSA still remains a secret.[2]

69.1.1.2 NSA PRISM in the United States.

In July 2013, the Electronic Privacy Information Center (EPIC) reported on its efforts to block the widespread collection of call detail records (metadata) about phone calls originating or terminating in the United States:

> EPIC has filed a petition with the U.S. Supreme Court, asking the Court to vacate an unlawful order by the Foreign Intelligence Surveillance Court that enables the NSA's collection of all domestic phone records. On April 25, the secret court ordered Verizon to turn over all "call

detail records" for calls made "wholly within the United States, including local telephone calls."

The FISC's order is based on Section 215 of the USA PATRIOT Act, which allows the court to compel the disclosure of business records that are "relevant to an authorized investigation." The Verizon order, however, requires the disclosure of all telephone records in the company's database.

"It is simply not possible that every phone record in the possession of a telecommunications firm could be relevant to an authorized investigation," EPIC states. "Such an interpretation of [the law] would render meaningless the qualifying phrases contained in the provision and eviscerate the purpose of the Act. "To define the scope of the records sought as 'everything' nullifies the relevance limitation in the statute," EPIC continues. "If law enforcement has 'everything,' there will always be some subset of 'everything' that is relevant to something."

The call detail records provided to the NSA, called "telephony metadata," contain an immense amount of sensitive personal information. The records identify the phone numbers of both parties on a call, the call's time and duration, and the geographic location of each phone number. When aggregated, such records can map out "the daily activities, interactions, personal and business relationships, religious and political affiliations, and other intimate details of millions of Americans," says EPIC.

EPIC's petition asks the Supreme Court to issue a "writ of mandamus" vacating the Verizon Order issued by the FISC. Mandamus is a command from a higher court to a lower court or government official, used when a lower court extends beyond the scope of its legal authority. EPIC writes, "Mandamus relief is warranted because the FISC exceeded its statutory jurisdiction when it ordered the production of millions of domestic telephone records that cannot plausibly be relevant to an authorized investigation."

EPIC brought the petition directly to the Supreme Court because no other court has jurisdiction to address the unlawful order. The Foreign Intelligence Surveillance Act (FISA) allows only the government and the recipient of a production order (i.e., Verizon) to appear before the FISC or the Court of Review. Additionally, the law limits these courts to hearing only certain types of appeals. EPIC, a Verizon customer whose call records are subject to the order, would not be able to obtain relief from these courts.[3]

69.1.1.3 Phone Hacking in the United Kingdom.

The Telegraph provides a detailed timeline of the events known as the U.K. phone-hacking scandal.[4] Key events include

- 2005 (November): The Royal Family complains about interception of private voicemail messages revealing a knee injury to Prince William.
- 2006 (August): "Detectives arrest the *News of the World*'s royal editor Clive Goodman and private investigator Glenn Mulcaire over allegations that they hacked into the mobile phones of members of the royal household."
- 2007 (January): "The *News of the World*'s royal affairs editor Clive Goodman is jailed for four months. Private investigator Glenn Mulcaire is given a six-month prison term. Goodman and Mulcaire admitted conspiring to intercept communications while Mulcaire also pleaded guilty to five other charges of intercepting voicemail messages."
- 2009 (July): "It emerges that *News of the World* reporters, with the knowledge of senior staff, illegally accessed messages from the mobile phones of celebrities and politicians while Coulson was editor from 2003 to 2007. It is also reported that News Group Newspapers, which publishes the *News of the World*, has paid out more than £1 million to settle cases that threatened to reveal evidence of its journalists' alleged involvement in phone hacking."

- 2011 (January): "British police open a new investigation into allegations of phone hacking at the tabloid called 'Operation Weeting' after actress Sienna Miller, MP George Galloway, and RMT union leader Bob Crow claim their phones were hacked."
- 2011 (February): "The Met Police release a statement saying officers have identified more potential victims of hacking while reviewing files relating to the original Goodman and Mulcaire case. They say they are urgently notifying people who had previously been told that police had 'little or no information' about them.

  "Former Deputy Prime Minister Lord Prescott, Labour MP Chris Bryant, ex-Scotland Yard commander Brian Paddick and journalist Brendan Montague, all alleged victims of phone hacking, win a High Court bid for a judicial review into the police inquiry. They believe their human rights were breached.

  "Lawyers for a football agent suing the *News of the World* claim Glenn Mulcaire passed information directly to the newsdesk rather than an individual reporter, Goodman. They say the desk could have been staffed by 'a number of journalists,' and suggest that this means knowledge of phone-hacking was more widespread than previously admitted. A judge rules that Mulcaire must provide information about whether other journalists at the *NoW* were involved in hacking. He had tried to claim he should be exempt from giving evidence for fear of incriminating himself."
- 2011 (March): "The BBC's Panorama reveals that in 2006, a then *News of the World* executive, Alex Marunchak, obtained e-mails belonging to an ex-British Army intelligence officer that had been hacked in to by a private detective."
- 2011 (April): "Former *News of the World* editor Ian Edmondson, chief reporter Neville Thurlbeck and senior journalist James Weatherup are arrested on suspicion of conspiring to intercept mobile phone messages. They are released on bail until September. The *News of the World* admits it had a role in phone hacking. The *News of the World* publishes apologies on both its website and newspaper. News International also announces it will set up a compensation scheme to deal with 'justifiable claims' fairly and efficiently. However, the publisher adds it will continue to contest cases 'that we believe are without merit or where we are not responsible.' "
- 2011 (May): "Former Deputy Prime Minister Lord Prescott, Labour MP Chris Bryant, ex-Scotland Yard commander Brian Paddick and journalist Brendan Montague, all alleged victims of phone hacking, win a High Court bid for a judicial review into the police inquiry. They believe their human rights were breached."
- In November 2012, Lord Justice Leveson released his report on the scandal.[5] Some of the key findings in the *Executive Summary* follow:
- "The evidence placed before the Inquiry has demonstrated, beyond any doubt, that there have been far too many occasions over the last decade and more (itself said to have been better than previous decades) when these responsibilities, on which the public so heavily rely, have simply been ignored. There have been too many times when, chasing the story, parts of the press have acted as if its own code, which it wrote, simply did not exist. This has caused real hardship and, on occasion, wreaked havoc with the lives of innocent people whose rights and liberties have been disdained. This is not just the famous but ordinary members of the public, caught up in events (many of them, truly tragic) far larger than they

could cope with but made much, much worse by press behaviour that, at times, can only be described as outrageous."

- "For many years, there have been complaints that certain parts of the press ride roughshod over others, both individuals and the public at large, without any justifiable public interest. Attempts to take them to task have not been successful. Promises follow other promises. Even changes made following the death of Diana, Princess of Wales, have hardly been enduring. Practices discovered by the Information Commissioner, during Operation Motorman, which led to the publication of two reports to Parliament,[6] revealed that large parts of the press had been engaged in a widespread trade in private and confidential information, apparently with little regard to the public interest. A private detective, Steve Whittamore, had certainly been engaged in wholesale criminal breaches of data protection legislation and, prima facie, journalists who engaged his services or used his products (and paid substantial sums for the privilege) must or should have appreciated that the information could not have been obtained lawfully."
- "When Clive Goodman, a journalist employed by the *News of the World* and Glenn Mulcaire, a private detective, were convicted of hacking into the telephone messages of members of the Royal Household and others, it was implicit during the course of the criminal prosecution that others must have been involved, whether knowingly or not, in using information that was the product of phone hacking. Most responsible corporate entities would be appalled that employees were or could be involved in the commission of crime in order to further their business. Not so at the *News of the World*. When the police had sought to execute a warrant, they were confronted and driven off by the staff at the newspaper. Cooperation, if provided, was minimal. The two that were sentenced to terms of imprisonment were paid very substantial sums as compensation for loss of employment when they were released."
- Then came exposure of the fact, albeit as long ago as March 2002, that the mobile phone of Milly Dowler had been hacked by someone at the *News of the World*. The information obtained had led the paper to publish false leads as a result of its misunderstanding of a message which had simply been left on the wrong phone in error. It was also believed that a message or messages had been deleted (thereby giving rise to a false moment of hope in her family). On 5 July 2011, these facts were reported by the *Guardian*. The outcry was immediate; two days later it was announced that the *News of the World* would close."
- "Phone hacking in itself, even if it were only in one title, would justify a reconsideration of the corporate governance surrounding the way in which newspapers operate and the regulatory regime that is required. Without making findings against anyone individually, the evidence drives me to conclude that this was far more than a covert, secret activity, known to nobody save one or two practitioners of the 'dark arts.' Yet it was illegal. And after the prosecution, at more than one title and more than one publisher, there was no in-depth look to examine who had been paid for what and why or to review compliance requirements."

69.1.2 Laws, Regulations, and Agreements. Across the emerging body of global privacy law, which varies substantially from jurisdiction to jurisdiction, general patterns are emerging. Virtually any piece of information related to an individual is defined as *personal information* (or *personally identifiable information,* PII*)* and, therefore, potentially subject to regulation with respect to its collection, storage, use,

and transfer. Different types of personal information, and different uses of personal information, are subject to different levels of regulation. Typically, the collection and use of medical and other highly sensitive data are most highly regulated; for example, data revealing the subject's race, religion, political affiliation, sexual orientation, or data relating to children.

What conduct is regulated? Privacy laws are structured on the essential principle that the individual person—a *data subject*—should retain some control over the information about himself or herself that is collected or used by a person or organization, whether a business, a nonprofit organization, or a government. Privacy laws impose obligations that are intended to prevent the collection or use of information in any manner that is inconsistent with the expectations of the data subject. In effect, the data subjects' control of their own personal information cannot be circumvented or denied, except with their consent. Privacy laws impose obligations on organizations that collect personal information, and empower the data subject and the state to enforce the data subject's rights under that agreement.

In addition, across different legal systems and, particularly in the United States, across different regulatory schemes governing different industries, privacy laws embrace comparable core principles, commonly referred to as *fair information practices*. These practices include the concept of control by the data subject but also may embrace rules that prohibit or limit certain data practices regardless of the consent of, or other involvement by, the data subject. The Office of the Privacy Commissioner of Australia summarizes these emerging principles as its *National Privacy Principles* (NPPs), as follows (quoting directly):

- NPP 1: collection. Describes what an organisation should do when collecting personal information, including what they can collect, collecting from third parties and, generally, what they should tell individuals about the collection.
- NPP 2: use and disclosure. Outlines how organisations may use and disclose individuals' personal information. If certain conditions are met, an organisation does not always need an individual's consent to use and disclose personal information. There are also rules about direct marketing.
- NPPs 3–4: information quality and security. An organisation must take steps to ensure the personal information it holds is accurate and up-to-date, and is kept secure from unauthorised use or access.
- NPP 5: openness. An organisation must have a policy on how it manages personal information, and make it available to anyone who asks for it.
- NPP 6: access and correction. Gives individuals a general right of access to their personal information, and the right to have that information corrected if it is inaccurate, incomplete or out-of-date.
- NPP 7: identifiers. Generally prevents an organisation from adopting an Australian Government identifier for an individual (eg Medicare numbers) as its own.
- NPP 8: anonymity. Where possible, organisations must give individuals the opportunity to do business with them without the individual having to identify themselves.
- NPP 9: transborder data flows. Outlines how organisations should protect personal information that they transfer outside Australia.
- NPP 10: sensitive information. Sensitive information includes information relating to health, racial or ethnic background, or criminal records. Higher standards apply to the handling of sensitive information.[7]

For chief information officers, security officers, risk managers, general counsels, systems designers, and consultants of any company charged with assuring compliance with privacy laws, new technologies and new business practices are raising new challenges for the enforcement of privacy rules. Regulatory controls that target information assets such as personal information are transforming how information systems are designed and managed. Technology executives are required to assume increased responsibility for legal compliance across divergent legal regimes; business executives and legal officers are required to acquire significant new understandings of the technologies deployed in information management and the business relationships through which personal information is collected, used, and transferred in support of the execution of business relationships. In some circumstances, information technology is both a source and a means of addressing privacy issues.

What makes privacy law important to all of these participants is not the imposition of yet another set of regulations requiring compliance. Instead, privacy law tests companies in their ability to truly integrate business, law, and technology into the coherent planning and execution of their information systems. As the world's economy continues its remarkable transformation into a global environment in which information in electronic form is a significant measure of economic wealth, those companies that master privacy as a competitive opportunity, rather than as a regulatory cost without economic benefit, will be best positioned to safeguard the data of their customers and avoid the growing concerns about security breaches and identity theft.

However, although attention is increasingly paid to overcoming inadequate protection of personal privacy, there is a corresponding increase in the attention paid to the policy issues raised by the impact of new privacy practices. A few examples will suffice. Law enforcement and national security agencies are concerned about how increased privacy protection, including encryption and anonymizer techniques, will weaken their ability to detect crime and terrorism. Medical researchers are concerned about their ability to conduct statistically valid medical research. At the same time, the absence of effective privacy safeguards has led to growing concerns about identity theft and security breaches, particularly in the United States.

69.1.3 Sources of Privacy Law. The legal structures governing personal information use have developed along different lines concerning public and private entities. Typically, governments and other public-sector entities are restrained from undue intrusion into citizens' private lives by constitutional mechanisms.[8] Citizens typically are provided access to government-held personal information by statutes, although there has been an increase in constitutionally mandated access rights to personal information as formerly oppressive regimes give way to more democratic governments. The converse is equally true.[9]

By contrast, restraints and obligations placed on private uses of personal information derive almost exclusively from statutory law. However, the recent adoption of the European Charter of Fundamental Rights requires that nation-states treat personal data protection as a basic human right. Such a requirement will stand as a prerequisite for any new potential member of the European Union.[10]

This development suggests a new trend: that data protection may be enshrined into constitutions or supranational law. Indeed, certain developing nations that have recently enacted new constitutions have created bases for their private-sector data protection laws in them. Also, the new member states of the European Union have all adopted data protection rules to comply with the EU Data Directive. As a result, any attempt to master what privacy law is all about will demand, regrettably perhaps, complex legal analysis.

69.2 EUROPEAN APPROACHES TO PRIVACY

69.2.1 History and Organization for Economic Cooperation and Development Principles. Interest in the right of privacy increased in the 1960s and 1970s with the advent of information technology and new data processing techniques for records concerning the individuals. The surveillance potential of computers prompted public demands for specific rules governing the collection and handling of personal information. The genesis of legislation in this area usually is traced to the first modern data protection law, enacted in the German state of Hesse in 1970. That enactment was followed by national laws with differing objectives and scope in Sweden (1973), the United States (1974), Germany (1977), and France (1978).[11]

Two highly influential international instruments emerged following these enactments. The Council of Europe (COE) passed the *Convention for the Protection of Individuals with regard to the Automatic Processing of Personal Information* (the *COE Convention*) in 1981.[12] In the same year, the Organization for Economic Cooperation and Development (OECD) issued *Guidelines Governing the Protection of Privacy and Transborder Data Flows of Personal information* (the *OECD Guidelines*).[13] Both instruments articulate specific rules regarding the handling of personal information. Both present a comparable set of principles for fair information use that have since become the framework for the data protection laws of dozens of countries. These principles recognize personal information as data to be afforded legal protection at every step of their life cycle of collection, processing, storage, use, transfer, and destruction.

To date (July 2013), 46 countries have adopted the COE Convention.[14] In addition to being relied on by OECD nations to create data protection laws, the OECD Guidelines have been relied on by other nations that are not OECD members.[15] The OECD has continued to be involved in privacy issues and, indeed, makes available a privacy policy statement generator on the privacy page of its Website.[16]

69.2.2 European Union Data Protection Directive 95/46/EC. In 1990, both to (1) ensure that member states protected "fundamental" privacy rights when processing personal information, regardless of the national citizenship of the data subjects, and (2) prevent member states from restricting the free flow of personal information within the European Union (EU), the European Commission proposed the EU Directive on Data Protection. On October 24, 1995, following years of discussion and debate, the Council and Parliament of the EU adopted the final version of Directive 95/46/EC.[17] The directive became effective on October 25, 1998.

The directive is not self-executing. Instead, it requires EU member states to pass national legislation that transpose its terms. It is important to realize that the national laws that have been enacted in each of the EU member states are not identical. Indeed, there are fairly significant differences among these laws. Thus, having an understanding of the directive by no means guarantees that a business understands all of the steps it must take to comply with each national law.

69.2.2.1 European Union Directive Requirements. The EU Directive applies to all commercial uses of personal information. The directive establishes limitations and obligations on *controllers* (entities directing the collection, use, or transfer of personal information) and *processors* (those utilizing personal information at another party's direction) with respect to the processing (i.e., virtually any collection, use, transfer, or destruction of personal information) of personal information (any

information that is readily traceable to a particular individual). The directive, like other privacy laws, is organized around certain key principles of information use. These can be summarized as follows:

- ***Notice.*** Data subjects, individuals to whom personal information relates, must be informed of:
  - The identity of the collector of the information;
  - The uses or purposes for which it is collected;
  - How they may exercise any available choices regarding its use;
  - Where and to whom it may be transferred; and
  - How data subjects may access information relating to themselves by an organization.
- ***Consent.*** Data subjects have the right, except under certain exceptions, not to have their personal information used without their consent, the right to opt out of having personal information used for direct marketing purposes, and the right not to have any "sensitive information" collected or used without express permission. Sensitive information includes information relating to race, religion, health, union membership, beliefs, and sex life.
- ***Consistency.*** Organizations may use personal information only in strict accordance with the terms of the notice given to the data subjects and any choices with respect to its use exercised by them.
- ***Access.*** Organizations must give data subjects access to the personal information held about them and must allow them to propose corrections to inaccurate information.
- ***Security.*** Organizations must provide adequate security, using both technical and other means, to protect the integrity and confidentiality of personal information. The sufficiency of such means is measured with respect to the state of the art.
- ***Onward Transfer.*** Personal information may not be transferred to a third party unless that third party has been contractually bound to use the data consistently with the notice given to data subjects, any choices they have exercised, and applicable law.
- ***Enforcement.*** The directive grants a private right of action to data subjects where organizations do not follow the law. In addition, each EU country has established a Data Protection Authority—a regulatory enforcement agency—that has the power to investigate complaints, levy fines, initiate criminal actions, and demand changes in information-handling practices of businesses.

69.2.2.2 International Data-Transfer Restrictions. While the directive governs much of the interaction between a business and its data subjects (i.e., its customers and employees), it also regulates how companies operating in multiple jurisdictions worldwide may move personal information within themselves and among their various affiliates or subsidiaries.

Transfers of personal information relating to EU-based data subjects to non-EU countries are *prohibited* unless the country of destination has *adequate legal protections for privacy*. Many countries, including the United States, are not considered by the EU to have adequate legal protection. Transfers to such countries may take place only with the consent of the data subject, or where other guarantees exist that personal

information will be treated according to the EU standard. EU regulators have the legal authority to prohibit transfers of personal information where such guarantees cannot be demonstrated. This requirement is problematic for many businesses and may require significant technological, organizational, and/or contractual changes.

One possible solution to the directive's transfer restrictions is the United States/European Union "Safe Harbor." This arrangement is discussed in Section 69.3.3.5.

69.2.2.3 State of Implementation. As of July 2013, the 28 member countries of the European Union, including the new members states, have passed legislation fully implementing the Privacy Directive.[18]

In January 2012, the EU received a plan for improving privacy protection:

> The European Commission has … proposed a comprehensive reform of the EU's 1995 data protection rules to strengthen online privacy rights and boost Europe's digital economy. Technological progress and globalisation have profoundly changed the way our data is collected, accessed, and used. In addition, the 27 EU Member States have implemented the 1995 rules differently, resulting in divergences in enforcement. A single law will do away with the current fragmentation and costly administrative burdens, leading to savings for businesses of around €2.3 billion a year. The initiative will help reinforce consumer confidence in online services, providing a much needed boost to growth, jobs, and innovation in Europe.[19]

According to a press release,

> Key changes in the reform include:
>
> - A single set of rules on data protection, valid across the EU. Unnecessary administrative requirements, such as notification requirements for companies, will be removed. This will save businesses around €2.3 billion a year.
> - Instead of the current obligation of all companies to notify all data protection activities to data protection supervisors—a requirement that has led to unnecessary paperwork and costs businesses €130 million per year, the Regulation provides for increased responsibility and accountability for those processing personal data.
> - For example, companies and organisations must notify the national supervisory authority of serious data breaches as soon as possible (if feasible within 24 hours).
> - Organisations will only have to deal with a single national data protection authority in the EU country where they have their main establishment. Likewise, people can refer to the data protection authority in their country, even when their data is processed by a company based outside the EU. Wherever consent is required for data to be processed, it is clarified that it has to be given explicitly, rather than assumed.
> - People will have easier access to their own data and be able to transfer personal data from one service provider to another more easily (right to data portability). This will improve competition among services.
> - A "right to be forgotten" will help people better manage data protection risks online: people will be able to delete their data if there are no legitimate grounds for retaining it.
> - EU rules must apply if personal data is handled abroad by companies that are active in the EU market and offer their services to EU citizens.
> - Independent national data protection authorities will be strengthened so they can better enforce the EU rules at home. They will be empowered to fine companies that violate EU data protection rules. This can lead to penalties of up to €1 million or up to 2% of the global annual turnover of a company.
> - A new Directive will apply general data protection principles and rules for police and judicial cooperation in criminal matters. The rules will apply to both domestic and cross-border transfers of data.[20]

Readers committed to following the Directive will want to keep tabs on developments as these proposals wend their way through the EU deliberative process.

69.2.3 Harmonization of Non–EU Countries to the EU Directive. Largely due to the EU Directive's prohibition on the transfer of personal information to countries with lesser degrees of legal protection, and the fear in many countries of the potential adverse economic impact that could result from the interruption of data flows from EU countries, a number of other countries have passed essentially identical national legislation so as to assure adequate levels of legal protection. These countries can be roughly divided into two categories: EU trading partners that wish to ensure that data flows will not be interrupted[21] and countries actively conforming their national legislation to the EU model in hopes of gaining membership to the EU.[22]

69.2.4 European Union Telecommunications Directive. In addition to the general Data Protection Directive, the EU has enacted a second directive, imposing specific privacy-protective obligations on communications service providers. The *Telecommunications Directive*[23] obliges carriers and service providers to ensure technologically the privacy of the communications carried on their systems, including email and other Internet communications. This directive also restricts access to billing information, places limitations on marketing, allows for per-line blocking of caller ID, and requires that information gathered during a communication be deleted on the termination of the communication.[24] A new proposal would replace the Telecommunications Directive with a broader version affecting all "electronic communications" in the European Union.[25]

69.2.5 Establishment of the European Data Protection Supervisor. In 2002, the European Parliament and the European Council established the office of the *European Data Protection Supervisor* (EDPS), described as follows:

> The EDPS is an independent supervisory authority devoted to protecting personal data and privacy and promoting good practice in the EU institutions and bodies. He does so by:
>
> - monitoring the EU administration's processing of personal data;
> - advising on policies and legislation that affect privacy; and
> - cooperating with similar authorities to ensure consistent data protection. …
>
> The supervisory task is to ensure that the EU institutions and bodies process personal data of EU staff and others lawfully. The EDPS oversees Regulation (EC) 45/2001 on data protection, which is based on two main principles:
>
> The responsible data controller needs to respect a number of obligations. For instance, personal data can only be processed for a specific and legitimate reason which must be stated when the data are collected.
>
> The person whose data are processed—the data subject—enjoys a number of enforceable rights. This includes, for instance, the right to be informed about the processing and the right to correct data.
>
> Every institution or body should have an internal Data Protection Officer. The DPO keeps a register of processing operations and notifies systems with specific risks to the EDPS. The EDPS prior checks whether or not those systems comply with data protection requirements. The EDPS also deals with complaints and conducts inquiries.[26]

In July 2013, the EDPS issued a report strongly criticizing the European Commission's plans for implementing computer-assisted border checks (*smart borders*). Peter Hustinx, the Director of the EDPS, was quoted as saying,

> In the absence of such a policy, the creation of yet another large-scale IT database to store massive amounts of personal information is a disproportionate response to a problem that other recently created systems may be able to help solve. It would be prudent both economically and practically to evaluate the existing systems at least to ensure consistency and best practice.[27]

On July 31, 2013, the positions of EDPS and Assistant Supervisor became vacant.[28] It was not known whether there was a relationship between the EDPS criticism of smart borders and the vacancy.

69.3 UNITED STATES

69.3.1 History, Common Law Torts. The concept of privacy violations as a tort is a twentieth-century phenomenon. The framers of the U.S. Constitution did not recognize an explicit right to privacy, perhaps because privacy was intrinsic to an agrarian society. But the necessity for privacy protection under law gradually became apparent with urbanization and industrialization in the late nineteenth and early twentieth centuries. The concept of a "right to be left alone" was posited by Charles Warren and Louis Brandeis in an influential 1890 article in the *Harvard Law Review*.[29] Today, virtually every state recognizes a common law or statutory right of privacy.[30]

Like most common law torts, privacy law matured in a piecemeal fashion at the state level because there was no federal law on point. In 1960, the *Restatement of Torts* defined four *subtorts* that are generally held to comprise the common law of privacy. These are[31]:

1. **Intrusion.** The unreasonable intrusion, by physical or other means, upon the seclusion of another, if the intrusion is substantial and would be highly offensive to a reasonable person.
2. **Revelation of private facts.** The unauthorized and unreasonable publicity of facts about a person's private life, if the matter publicized is of a kind that would be highly offensive to a reasonable person, and which is not of legitimate concern to the public. These facts must be communicated to more than a single person or a small group in order to trigger this tort.
3. **False light.** Publicity that unreasonably places the other in a false light before the public. Essentially, a photo or story cannot be used if that photo or story conveys a false impression of someone. Note that this tort can be used if the person is cast in either a negative or *positive* false light.
4. **Misappropriation.** The unauthorized appropriation of another's name or likeness for benefit or gain. Misappropriation actions are most commonly brought by celebrities whose image is used without authorization in connection with advertising or promotion.

69.3.2 Public Sector. In the United States, government traditionally has been restrained from undue intrusion into the private lives of citizens by various provisions of the U.S. Constitution and similar provisions in the constitutions of many of the states of the Union. The Fourth Amendment to the U.S. Constitution, for instance, prohibits

"unreasonable search and seizure" by the government. Fourth Amendment protection is also made applicable to the states by operation of the Fourteenth Amendment to the U.S. Constitution, which provides that "[n]o state shall make or enforce any law which shall abridge the privileges or immunities of citizens of the United States." Yet, except in certain state constitutions, there is no explicit right of privacy. Thus, many of the rules regarding how government may collect, maintain, and use personal information relating to citizens has been defined by case law and statutes. In the public sector, state governments may have greater privacy protection obligations to their employees than private sector businesses in the same states have to their employees.

69.3.2.1 Privacy Act of 1974 (Including the Freedom of Information Act). The most comprehensive privacy legislation in the United States, the Privacy Act of 1974,[32] placed limitations on how the Federal Government may use and transfer personal information, and gave individuals a right, subject to certain exceptions, to know what information is held about them by the Federal Government.[33] The Privacy Act incorporates the Freedom of Information Act (FOIA), which allows access to Federal Government records, except in certain situations.

The Act is intended to:

- Permit individuals to determine what records pertaining to them are collected, maintained, used, or disseminated by federal agencies
- Permit individuals to forbid records that are obtained by an agency for one purpose to be used for another purpose without consent
- Permit individuals to obtain access to agency records pertaining to them and to correct or amend such records as appropriate
- Ensure that agencies collect, maintain, and use personal information in a manner that ensures that the information is current, adequate, relevant, and not excessive for its intended use
- Create a private right of action for individuals whose personal information is not used in accordance with the act[34]

All provisions of the Act are subject to certain exceptions, such as criminal investigations or national security concerns. In addition, the Act protects each individual, so access to certain information can be limited where providing such information would compromise another individual's privacy.

69.3.2.2 Electronic Communications Privacy Act of 1986. The Electronic Communications Privacy Act (ECPA),[35] which amended the federal wiretap law of 1968, generally prohibits unauthorized and intentional "interception" of wire, oral, and electronic communications during the transmission phase, and unauthorized "accessing" of electronically stored wire or electronic communications. It is applicable in many connections. For example, it establishes standards for capturing telephone numbers through the use of pen registers and trap-and-trace devices. Pen registers record telephone numbers of outgoing calls; trap-and-trace devices record telephone numbers from which incoming calls originate, much like common caller ID systems. Although telephone numbers are not protected by the Fourth Amendment,[36] the ECPA requires law enforcement agencies to obtain court orders to install and use these devices. Rather than the strict probable cause showing the necessity for wiretaps, pen register orders require only certification from a law enforcement officer that "the information likely

to be obtained is relevant to an ongoing criminal investigation."[37] It is also relevant to workplace invasions of privacy by electronic means (see Section 69.3.3.7).

69.3.2.3 Right to Financial Privacy Act of 1978. Under the Right to Financial Privacy Act,[38] the federal government is prohibited from obtaining the financial records of an individual, or information contained therein, from a financial institution without giving notice to the individual. Where the disclosure is sought via a subpoena, the required disclosure normally must be made within 90 days after the subpoena is served on the financial institution. Where the disclosure is sought by all other means, the required notice normally must be given prior to, or contemporaneously with, the disclosure request.

69.3.2.4 Driver's Privacy Protection Act. The Driver's Privacy Protection Act[39] marked the first time the U.S. Congress passed a law limiting how state governments may use personal information. Specifically, the Act prohibits disclosures of personal information obtained in connection with a motor vehicle record by state governments, except in order to facilitate legitimate government activities, facilitate recalls, and the like. This Act provides civil and criminal penalties as well as a private right of action. The Act survived a challenge to its constitutionality, brought by the state of South Carolina, by a 9-to-0 vote in the U.S. Supreme Court.[40]

69.3.2.5 Law Enforcement and National Security Surveillance. As technology has developed, crime and criminal activities have become increasingly sophisticated. Yet the same developments of technology have greatly facilitated information gathering and surveillance by law enforcement and national security organizations.

Law enforcement agencies have used technology to investigate crimes such as money laundering, narcotics trafficking, insider trading, and fraud. The technology used can be divided roughly into two categories: monitoring and surveillance. Monitoring technology is used to search for indicia of criminal behavior, without focusing on the activities of particular individuals. Examples of monitoring technologies used to detect crime online include:

- **Packet sniffers.** Programs that monitor all traffic flowing through a network, looking for keywords, phrases, email addresses, and the like. When a packet sniffer finds the type of data it is looking for, it saves the communication for later review.
- **"Black boxes."** Hardware installed by law enforcement agencies at Internet service providers (ISPs) in order to monitor communications or Web traffic. Typically, but not always, black boxes use packet-sniffing technology. In the United States, the Federal Bureau of Investigation (FBI) promoted the use of a black box system installed at ISPs called "carnivore" later renamed DCS1000[41]; currently the program is under review. In February 2001, the FBI renamed its software DCS1000.

Surveillance technology, by contrast, is used to eavesdrop on the communications or transactions of particular individuals. Monitoring technologies often are used or usable for this purpose. In addition, audio listening devices (bugs) and closed-circuit television (CCTV) are increasingly in use.[42]

National security agencies worldwide conduct monitoring and surveillance of communications, to varying degrees. The United States, United Kingdom, New Zealand,

Australia, and Canada are reputed to have an integrated global surveillance system known as Echelon. Although details of the Echelon system's capabilities are closely guarded state secrets, some reports indicate that the system is able to intercept virtually any electronic communication on earth.[43] Telecommunications operators in the United States should be aware that the Communications Assistance for Law Enforcement Act of 1994 (CALEA)[44] mandates certain technical standards for telecommunications networks in order to facilitate law enforcement interception of communications. Recent court decisions may, however, ultimately limit CALEA's application.[45] In 2005 the FCC's First Report and Order ruled that ISPs are "telecommunications carriers" under CALEA. This was challenged and upheld in the D.C. Circuit Court in 2006.

In 2001, the Council of Europe is developed, in consultation with the U.S. Department of Justice, an international Convention on Cyber-Crime which entered into force in 2004.[46] The Convention is controversial. As of September 2013, 40 countries have ratified the convention, and 11 are signatories but have not ratified the treaty.[47] Privacy advocates have also criticized this convention for being unduly intrusive.[48]

69.3.2.6 Foreign Intelligence Surveillance Act and Court. The Foreign Intelligence Surveillance Court (FISC) was established in 1978 as part of the Foreign Intelligence Surveillance Act (FISA). The Federal Judicial Center of the U.S. Government describes the origin and functioning of FISC and FISA as follows:

> Congress in 1978 established the Foreign Intelligence Surveillance Court as a special court and authorized the Chief Justice of the United States to designate seven federal district court judges to review applications for warrants related to national security investigations. Judges serve for staggered, non-renewable terms of no more than seven years, and until 2001 were drawn from different judicial circuits. The provisions for the court were part of the Foreign Intelligence Surveillance Act (92 Stat. 1783), which required the government, before it commenced certain kinds of intelligence gathering operations within the United States, to obtain a judicial warrant similar to that required in criminal investigations. The legislation was a response to a report of the Senate Select Committee to Study Governmental Operations with Respect to Intelligence Activities (the "Church Committee"), which detailed allegations of executive branch abuses of its authority to conduct domestic electronic surveillance in the interest of national security. Congress also was responding to the Supreme Court's suggestion in a 1972 case that under the Fourth Amendment some kind of judicial warrant might be required to conduct national security related investigations.
>
> Warrant applications under the Foreign Intelligence Surveillance Act are drafted by attorneys in the General Counsel's Office at the National Security Agency at the request of an officer of one of the federal intelligence agencies. Each application must contain the Attorney General's certification that the target of the proposed surveillance is either a "foreign power" or "the agent of a foreign power" and, in the case of a U.S. citizen or resident alien, that the target may be involved in the commission of a crime.
>
> The judges of the Foreign Intelligence Surveillance Court travel to Washington, D.C., to hear warrant applications on a rotating basis. To ensure that the court can convene on short notice, at least one of the judges is required to be a member of the U.S. District Court for the District of Columbia. The act of 1978 also established a Foreign Intelligence Surveillance Court of Review, presided over by three district or appeals court judges designated by the Chief Justice, to review, at the government's request, the decisions of the Foreign Intelligence Surveillance Court. Because of the almost perfect record of the Department of Justice in obtaining the surveillance warrants and other powers it requested from the Foreign Intelligence Surveillance Court, the review court had no occasion to meet until 2002. The USA Patriot Act of 2001 (115 Stat. 272) expanded the time periods for which the Foreign Intelligence Surveillance Court can authorize surveillance and increased the number of judges serving the court from seven to eleven. The eleven judges must be drawn from at least seven judicial circuits, and no fewer than three are to reside within twenty miles of the District of Columbia.[49]

Despite criticism by civil-liberties proponents, additional legislation was passed in 2007 and 2008 to strengthen the Executive Branch's powers to order wide-reaching surveillance without prior approval of the FISC.[50]

69.3.2.7 USA PATRIOT Act. The *Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism* (USA PATRIOT) Act, whose name was crafted to result in the positive-sounding acronym, was passed with minimal legislative review in October 2001 as a reaction to the terrorist attacks of September 11, 2001.[51]

The Electronic Frontier Foundation, a civil-liberties organization, summarizes its perception of the threats in the USA PATRIOT Act to online civil liberties as follows:

- The law dramatically expands the ability of states and the Federal Government to conduct surveillance of American citizens. The Government can monitor an individual's web surfing records, use roving wiretaps to monitor phone calls made by individuals "proximate" to the primary person being tapped, access Internet Service Provider records, and monitor the private records of people involved in legitimate protests.
- PATRIOT is not limited to terrorism. The Government can add samples to DNA databases for individuals convicted of "any crime of violence." Government spying on suspected computer trespassers (not just terrorist suspects) requires no court order. Wiretaps are now allowed for any suspected violation of the Computer Fraud and Abuse Act, offering possibilities for Government spying on any computer user.
- Foreign and domestic intelligence agencies can more easily spy on Americans. Powers under the existing Foreign Intelligence Surveillance Act (FISA) have been broadened to allow for increased surveillance opportunities. FISA standards are lower than the constitutional standard applied by the courts in regular investigations. PATRIOT partially repeals legislation enacted in the 1970s that prohibited pervasive surveillance of Americans.
- PATRIOT eliminates Government accountability. While PATRIOT freely eliminates privacy rights for individual Americans, it creates more secrecy for Government activities, making it extremely difficult to know about actions the Government is taking.
- PATRIOT authorizes the use of "sneak and peek" search warrants in connection with any federal crime, including misdemeanors. A "sneak and peek" warrant authorizes law enforcement officers to enter private premises without the occupant's permission or knowledge and without informing the occupant that such a search was conducted.
- The Department of Justice, with little input from Congress and the American people, is developing follow-on legislation—the Domestic Security Enhancement Act (nicknamed Patriot II)—which would greatly expand PATRIOT's already sweeping powers.[52]

69.3.3 Private Sector. Compared to other nations, the United States has been fairly slow to pass privacy laws affecting the private sector, preferring instead to rely on industry self-regulation (see Section 69.4.1). Nonetheless, the development of the Internet and electronic commerce, the integration of the financial services industries, and the integration of healthcare services have all spawned new private-sector

legislation. Thus far, this legislation has concentrated on particular business sectors rather than regulating uses of personal information generally and has generally emphasized online business practices.

69.3.3.1 Financial Services Sector. The Gramm-Leach-Bliley Act (GLBA), enacted in 1999, requires every *financial institution* to:

- Protect the security and confidentiality of its customers' nonpublic personal information,
- Disclose its privacy policies to consumers, and
- Provide consumers with an opportunity to direct that the institution not share their nonpublic personal information with unaffiliated third parties.[53]

Violations of the privacy provisions of GLBA are subject to enforcement by various federal regulators; the specific sanctions that may be imposed vary according to the regulatory agency that has jurisdiction in a particular case.

Eight federal agencies have issued regulations enacting the privacy provisions of GLBA.[54] The Federal Trade Commission (FTC), among others, has acted aggressively to apply GLBA as broadly as possible. Specifically, the FTC defines *financial institution* by reference to a variety of other regulations:

> Under the Federal Trade Commission's Privacy Rule, a financial institution means "any institution the business of which is engaging in financial activities as described in § 4(k) of the Bank Holding Company Act of 1956 [12 U.S.C. § 1843(k)]." See 16 C.F.R. § 313.3(k)(1). (1) Further, you are not a financial institution unless you are significantly engaged in financial activities. Id. Moreover, the Commission's Rule covers only financial activities as of the date of enactment of the GLB Act. See 16 C.F.R. § 18(a)(2). Thus, financial activities added by regulation or order by the Federal Reserve Board or the Secretary of the Treasury after November 12, 1999, are not automatically covered by the Commission's Rule until the Commission so determines.[55]

This has the effect of causing a number of businesses not traditionally thought of as financial institutions, such as career counselors,[56] real estate settlement services, and property appraisers,[57] to be subject to GLBA's privacy regulations.

Another example of the broad application of these regulations is how *nonpublic personal information* is defined. In the case of certain agencies, it is defined as *any* personal information given by a consumer to an institution, whether the information is otherwise publicly available or not.

The regulations went into effect on July 1, 2001. All affected companies were required to have systems in place to comply with the regulations by that date. A list of the critical aspects of the GLBA regulations follows.

- *Customers*,[58] and in some instances *consumers*,[59] must receive initial and annual notices of how their personal information will be used by the financial institution and to which nonaffiliated third parties, other than service providers, it may be transferred.
- Where personal information is to be shared with nonaffiliated third parties other than service providers, the customer or consumer must be given an opportunity to *opt out* of having the personal information transferred.

- Appropriate security controls must be maintained, and specified procedures must be followed, for the adoption and testing of information-security protocols.
- Sharing of personal information among affiliates remains unaffected.

Some of the effects of GLBA that businesses should be aware of include:

- Systems must be created to capture a reliable record of a customer's opt-out (where applicable) at any time and to adopt controls to prevent that customer's personal information from being transmitted outside of a company or its affiliates.
- Systems must be created to deliver the appropriate notices to customers.
- Credit bureaus are no longer legally able to sell *credit header* information (i.e., as defined by the FTC in 1997, the identifying features of a credit report—name, aliases, date of birth, Social Security number, addresses current and prior, and current telephone number).
- Banks and other financial institutions are limited in their ability to sell or transmit lists of customers.
- Banks and other financial institutions must draft and execute, in most cases, detailed written information security policies.[60]
- Contract provisions governing the security of information transferred from financial institutions to other entities should be reviewed and may need to be amended.
- Where a particular company may be subject to GLBA regulation by more than one federal agency, particular care should be used to establish any differences in the different agencies' regulations and any provisions that may defer to a different agency's rules in such a situation.

There is an extensive series of practical questions and detailed answers about the Privacy Regulation on the FTC Website.[61] See Chapter 64 in this *Handbook* for a more extensive discussion of GLBA implications for security management.

69.3.3.2 Children's Online Privacy Protection Act. The Children's Online Privacy Protection Act of 1998 (COPPA) prohibits the collection, use, or disclosure of personally identifiable information from and about children under the age of 13 on the Internet without prior *verifiable parental consent*.[62]

The law firm of Kelly/Warner published an excellent summary of the COPPA provisions, including the following on the "Basic Provisions of COPPA:"

Who Must Abide by the COPPA Standards?

1. Anyone who operates an online service or commercial website that is attractive to children below the age of 13. Even if you don't mean to target kids, if they come to your site, you're responsible for knowing and following COPPA standards.
2. Anyone operating a website designed for general public consumption that collects information about users.
3. Anyone who develops plugins or add-ons that can be used on a site visited by minors.

Factors evaluated when determining if a website or software falls under COPPA's purview:

- Subject matter;
- Language;
- Advertising on the site, coupled with the intended target;

- Audio or visual content including animated characters;
- Age of the models;
- Information regarding the age of the intended or actual audience;
- Any features that may be child-oriented.

In addition, when determining who is and who isn't an "operator" under COPPA definitions, the FTC considers the quality of the information being collected. For example:

- Who owns and controls the information being collected?
- Who pays for the information and maintains it?
- What pre-existing contractual information exists? Is connected with that information?
- What is the role of the website under investigation in collecting and maintaining the information under review?

What Personal Information Are We Talking About When It Comes To COPPA?
COPPA's rules define what collected information is considered actionable under the law. Specifically:

- Full Name
- Home Address
- Email Address
- Telephone Number

Any other information that could be used to contact or identify a child is also covered, like interests and hobbies. Also covered under COPPA is digital info collected through Internet cookies or other tracking technologies, when are attached to personally identifying information. …

Basic Provisions of COPPA:
Under the Children's Online Privacy Protection Act a website's Privacy Notice must:

- include a link to the operation's information practices for each area where a child's personal information is collected.
- be posted prominently and clearly, using a font or text color distinctively different than the other text.
- be clearly stated without any confusing terminology or unrelated material that might distract from the intended message.
- sit on a contrasting background (i.e., dark text on a light background for easy reading).

Parental Notice Under COPPA
Privacy policies on websites should include a parental notice telling parents what, why and how information is being collected, in addition to what is done with the data. It must inform parents that the information can only be collected with verifiable parental consent. It should also instruct parents on how to submit consent as well as complaints.

Under COPPA, the terms or privacy notice should also include:

- The names of all website operators and the contact information of at least one;
- The information is being collected, with an explanation of how it's being collected;
- Explanation of how the collected data will be used;
- Third-party info: who they are, their business and how they will use the info;
- Clear notice on how parents can disallow third party access to the information;
- Instructions on how to get data deleted and the process of how to disallow further data collection.[63]

Obtaining verifiable parental consent has proven technologically challenging because of the ease with which children can deceptively click on confirmation of adult

status on Website pages attempting to limit access by minors. Even attempts relating to provision of credit card numbers have proven ineffective in barring some children from accessing sites. See Chapter 72 in this *Handbook* for a discussion of parental controls on access to the Internet and the Web.

Violations of COPPA are enforced as *unfair* or *deceptive* trade practices by the Federal Trade Commission.[64] Website operators should refer to the FTC's informational Website[65] for more information.

69.3.3.3 Health Insurance Portability and Accountability Act. The Health Insurance Portability and Accountability Act of 1996 (HIPAA) mandated the establishment of standards to protect the privacy and confidentiality of individually identifiable health information.[66] Under HIPAA, the Department of Health and Human Services has adopted rules to implement the legislative mandate.[67] Both civil and criminal penalties can be assessed for violations.

The HIPAA regulations present significant challenges to "covered entities." A number of the key concepts embodied in the regulations are not clearly defined. The obligations created by the regulations required significant changes in the systems infrastructure and business processes of many affected entities. HIPAA dramatically affects any business dealing with personal information gathered in a healthcare setting. Compliance with the HIPAA privacy rules became mandatory for covered entities in April 2003, while the compliance deadline for small health plans was April 2004. Extensive interpretive guidelines were published by Health & Human Services (HHS) in 2009.[68]

The key features of the HIPAA privacy regulations include:

- Providers and health plans are required to give patients a clear written explanation of how they will use, keep, and disclose information.
- Covered entities must provide patients with access to their health records.
- Permitted uses and disclosures of personally identifiable health information must, subject to certain exceptions, be limited to the "minimum amount necessary to accomplish the purpose for which the information is used or disclosed."[69]
- Disclosure logs must be kept that record each entity to which personally identifiable health information has been disclosed, and individuals must have a right to access these logs.
- A provider or payer generally is not able to condition treatment, payment, or coverage on a patient's agreement to the disclosure of health information for other purposes.
- Federal criminal penalties will apply to obtaining health information under false pretenses, and knowing and improper disclosures of health information, with penalties set higher if the violation is done for "monetary gain."
- States are not preempted from further regulation of health information, so long as an affected entity could comply with both the federal and the state requirements.
- Violations of the regulations will be subject to substantial fines.

In 2005, the *Patient Safety and Quality Improvement Act* (PSQIA) established

> . . . a voluntary reporting system designed to enhance the data available to assess and resolve patient safety and health care quality issues. To encourage the reporting and analysis of medical errors, PSQIA provides Federal privilege and confidentiality protections for patient

> safety information, called *patient safety work product*. PSQIA authorizes HHS to impose civil money penalties for violations of patient safety confidentiality. PSQIA also authorizes the Agency for Healthcare Research and Quality (AHRQ) to list patient safety organizations (PSOs). PSOs are the external experts that collect and review patient safety information. …
>
> The Patient Safety Rule implements select provisions of PSQIA.
>
> Subpart C of the Patient Safety Rule establishes the confidentiality provisions and disclosure permissions for patient safety work product and the enforcement procedures for violations of confidentiality pursuant to section 922 of the statute. OCR [Office for Civil Rights[70]] enforces these confidentiality protections.[71]

The HIPAA regulations are complex with respect to privacy and electronic data transactions.[72] The HHS provides an extensive range of *Frequently Asked Questions* that provide guidance for implementation.[73]

For more information on medical records security, see Chapter 71 in this *Handbook*.

69.3.3.4 Cable and Video Acts. Although the United States has been very hesitant to pass private-sector legislation regulating how businesses use personal information, it has not altogether shied away from doing so. Two important exceptions to that general policy stance are in the form of federal legislation: the *Cable Communications Policy Act*[74] and the *Video Privacy Protection Act*.[75]

The Cable Communications Policy Act of 1984, section 551, governs the *protection of subscriber privacy*. This section requires that cable television companies give their customers initial and annual notice of what personal information the companies collect, to whom it is distributed, how customers may access it, and what other limitations the Act places on the companies. Cable companies are prohibited from using their cable systems to collect personal information without prior consent of customers. Companies are also prohibited from retransferring any personal information without customers' prior consent. Companies may retransfer names, addresses, and telephone numbers to other companies, but only after offering customers the chance to opt out of such disclosures. Law enforcement may obtain personal information from cable companies by court order. The Act creates a private right of action for customers, which includes the ability to recover punitive damages and attorneys' fees from companies that violate it. An ongoing question concerns the application of these privacy provisions to new services offered by cable companies.

The Video Privacy Protection Act of 1988 prohibits the transfer of videotape rental and sales records except for certain narrow exceptions, including permission of the customer. Transfers of such records may be made to law enforcement agencies enforcing a warrant. Customers may bring civil suits against violators of this Act. A case started in 2007 involving the social network service Facebook asserted that Blockbuster violated the Act when it disclosed video purchases to Facebook subscribers using that service's *Beacon* program.[76]

69.3.3.5 Privacy of Activity in Libraries and Bookshops. In sharp contrast with the privacy provisions of privacy protections for cable television and video rentals and sales discussed in Section 69.3.3.4, the USA PATRIOT Act (see Section 69.3.2.7) specifically allows intelligence and law enforcement agencies to monitor the use of libraries and booksellers by patrons. Specifically, librarians are required to submit detailed borrowing records on demand and may not reveal that these records were surrendered.

> The Patriot Act permits federal agents to secretly obtain information from booksellers and librarians about customers' and patrons' reading, internet, and book-buying habits, merely by alleging that the records are relevant to an anti-terrorism investigation. The act prohibits librarians and booksellers from revealing these requests, so they cannot be challenged in court. … A University of Illinois study . . . concludes that federal agents have sought records from about 220 libraries nationwide since September 2001 . . . [as of January 2002]. The Justice Department refuses to say how many times it has invoked this Patriot Act provision . . . But Assistant Attorney General Daniel Bryant says that people who borrow or buy books surrender their right of privacy. … Some libraries and bookstores unhappy with the law begin to fight back in a number of ways. Some libraries have posted signs warning that the government may be monitoring their users' reading habits. … Thousands of libraries are destroying records so agents have nothing to seize. … Many librarians polled say they would break the law and deny orders to disclose reading records.[77]

The American Library Association explains their perspective on why the provisions matter to the public:

> Libraries are key sources of information on all kinds of subjects and from all perspectives for their communities. Reading has always been one of our greatest freedoms. The written word is the natural medium for the new idea and the untried voice from which come the original contributions to social growth.
>
> Libraries provide a place to exercise intellectual freedom: a free and open exchange of knowledge and information where individuals may exercise freedom of inquiry as well as a right to privacy in regards to information they seek. Privacy is essential to the exercise of free speech, free thought, and free association. In a library, the subject of users' interests should not be examined or scrutinized by others.
>
> The ALA believes certain sections of the USA PATRIOT Act endanger constitutional rights and privacy rights of library users. Libraries cooperate with law enforcement when presented with a lawful court order to obtain specific information about specific patrons; however, the library profession is concerned some provisions in the USA PATRIOT Act go beyond the traditional methods of seeking information from libraries.[78]

69.3.3.6 United States/European Union Safe Harbor. The EU Directive restricts the export of personal information to only those nations with laws offering "adequate" privacy protection.[79] Since, as of 1998, no official body in the EU had determined U.S. law to be "adequate," ongoing data transfers from Europe to the United States across a wide number of significant industries (airlines, entertainment, credit cards, retailing) were in jeopardy of disruption.

In April 1998, United States and European Commission officials entered into negotiations to create a "Safe Harbor" for American companies desiring to import from Europe personal information relating to European citizens, by which American companies (and their trading counterparts) would not be required to prove the adequacy of American law as a condition to the execution of data transfers. Difficult negotiations ultimately produced, in July 2000, an arrangement under which transfers of personal information from the European Union to the United States will be allowed for those companies willing to comply with the "Safe Harbor" principles.[80] In substance, the principles are similar to the requirements of the directive. Under those principles, in the absence of committing to the jurisdiction of European officials, an American company importing personal information on European citizens must:

- Comply with the Safe Harbor principles, according to the guidance offered in the FAQs (see www.export.gov/safeHarbor for more details).

- Self-certify, on a voluntary basis, with the U.S. Department of Commerce its adherence to principles for the protection of personal information that are consistent with European legal principles, which adherence must be publicly announced.
- Provide for independent third-party verification that the company adheres to the Safe Harbor principles. This step can be accomplished by joining a self-regulatory privacy program that adheres to the Safe Harbor. Joining such an organization can, by itself, also qualify the company for the Safe Harbor. Examples such as TRUSTe or BBBOnline have been cited as representative of the types of programs contemplated; in effect, the U.S. and EU governments are encouraging self-regulation, deferring to nongovernment programs to enforce suitable standards on participating companies.[81]

Companies subject to explicit federal regulations regarding the privacy of personal information (e.g., HIPAA)[82] may in the future be eligible to self-certify their compliance with the applicable laws and gain the Safe Harbor recognition that would permit their transactions with European data to occur.[83] In the meantime, companies must be subject to the jurisdiction of the Federal Trade Commission for unfair or deceptive trade practices, under the U.S. Federal Trade Commission Act, or subject to the jurisdiction of the U.S. Department of Transportation in order to be eligible to join the Safe Harbor.[84] Failure of a company that has joined the Safe Harbor to abide by its stated privacy policy is actionable as an unfair or deceptive trade practice.[85] In addition, companies would be subject to sanctions under other applicable laws.[86]

Currently, European officials in those nations that have enacted laws responsive to the EU Directive are authorized to enforce those laws against U.S. companies, or their trading partners, involved in exporting personal information. To date, few conspicuous actions have been taken (largely as a result of an informal "standstill" agreed on during the pendency of the Safe Harbor negotiations).[87] As of July 2013, there were 4,091 organizations on the Safe Harbor List. However, 944 organizations were designated "not current" with respect to certification status.

The "US–EU Safe Harbor" pages at export.gov provide extensive resources for helping organizations comply with these agreements.[88]

69.3.3.7 Workplace Privacy. Software has been developed, and is increasingly used, to monitor employee email, voicemail, and Web usage. Some software can record every keystroke made by an employee,[89] or allow an employer to see what Web page the employee has on the screen in real time.

The area of workplace privacy presents an especially difficult Hobson's choice for employers. Monitoring gives rise to assertions of invasion of privacy, and the failure to monitor gives rise to claims of negligence in failing to detect and act on claims of maintaining a discriminatory work environment.

In the United States, the common law is fairly well established that because employers own the computer equipment and networks used by their employees, and particularly where employees have been given explicit notice that their Web usage or email may be subject to monitoring, employers may monitor such computer usage.[90]

The Electronic Communications Privacy Act (ECPA) (see Section 69.3.2.2) is the only federal statute relevant to claims of workplace invasions of privacy by electronic means. ECPA prohibits unauthorized and intentional "interception" of wire, oral, and electronic communications during the transmission phase and unauthorized

"accessing" of electronically stored wire or electronic communications. An email is an "electronic communication."[91]

ECPA contains two important exceptions. The first is the "system provider" exception. There is considerable debate whether an employer who provides a company system for email and/or Internet access is covered by this exception. The second is the "consent" exception. Under this exception, an employer may intercept electronic communications if the prior consent of *one* of the parties to the communication has been obtained.[92] Hence, an employer needs the implied or express consent of *only* the employee. An employee probably will be considered to have given consent if he or she continues to use a communications system after having been advised of the employer's monitoring policy, although evidence of explicit consent through a signed writing, a returned email, or other technique is clearly preferable. For this reason, many companies require that employees agree to the company monitoring policy in writing as a condition of employment.

While voicemail is subject to employer monitoring by the same rules that apply to email, live employee telephone calls, even when made from company phones on company time, are subject to different rules. ECPA's prohibition on the "interception" of telephone calls during the transmission phase requires that employers not monitor employees' *non–work-related* telephone calls. Although it is awkward, the rule requires employers to cease monitoring of employees' telephone calls on the first indication that the call is made for personal purposes.[93]

In Florida and Maryland, state law requires that *both* parties to a communication consent to the monitoring of that communication.[94] Obtaining only an employee's consent therefore is unlikely to be legally sufficient to allow such monitoring in these states. In addition, other state laws governing the monitoring and interception of communications may apply. As of the publication deadline date of this chapter, however, there were no laws at either the state or the federal level prohibiting workplace email or Web usage monitoring in the United States.

In the European Union, all workplace monitoring is considered to be the collection of personal information and is subject to the terms of the Data Protection Directive, discussed earlier.

For additional discussions of employee management policies, see Chapter 45 in this *Handbook.*

69.3.3.8 Anonymity of Internet Postings and Defamation. Internet chat rooms, forums, and message boards have been used increasingly by disgruntled employees, customers, competitors, and others to attempt to damage or artificially elevate the stature of companies in the marketplace. Postings may include personal attacks; assertions about a company's business practices, the quality of its goods and services, or the worth of its stock. Such postings may be highly offensive to particular individuals, may result in intangible changes of perception about a company in the marketplace, or may result in a very real and measurable impact on the company's stock price. Understandably, companies that find themselves subject to such postings may wish to take action.

Options that a company may wish to consider when faced with negative or harmful Internet postings follow:

- **Do nothing.** Posters of negative messages may feed on the attention created by a public response. If ignored, they may become disinterested in pursuing the activity.

- **Utilize investigators to identify the poster.** If the poster's identity is known and the person is contacted by the company's counsel, the poster may cease the offending activity.
- **Contact law enforcement.** This option is appropriate where postings seem to have been intended to manipulate stock prices or where messages threaten or suggest harm to individuals or property.
- **File suit against "John Doe" and subpoena the relevant ISP to discover the identity of the poster.** This option is not as straightforward as it may seem. Many "John Doe" defendants have vigorously challenged such subpoenas,[95] and a recent court decision ruled that an ISP did not have to reveal the poster's identity.[96]

For a more detailed review of anonymity in cyberspace, see Chapter 70 in this *Handbook*.

69.3.3.9 Online Monitoring Technology. Much of the public concern surrounding privacy has derived from the use of various technologies to monitor Internet usage and compile information about users, often without clear notice to consumers. This section gives a brief description of the most prevalent technologies currently in use for online information gathering.

Cookies are small text files that are placed by a Web server onto the hard drive of an Internet user when that user visits a Website. Cookies allow Web servers to recognize computers bearing their cookies when users revisit Websites. Websites can use cookies to monitor what pages users view, what sites they come from to the Website, and where they go when they leave the site; cookies also can be used to trigger stored user settings and preferences at the Website. The use of cookies is almost ubiquitous on commercial Websites. Users have the ability, through browser software, to view and delete cookies or to set their browsers to reject them or warn them each time a cookie is placed.

The Electronic Privacy Information Center (EPIC) has a long history of what it defines as abuse of cookies:

> EPIC filed a complaint . . . with the Federal Trade Commission on February 10, 2000, concerning the information collection practices of DoubleClick Inc., a leading Internet advertising firm, and its business partners. The complaint alleges that DoubleClick is unlawfully tracking the online activities of Internet users (through the placement of cookies) and combining surfing records with detailed personal profiles contained in a national marketing database. EPIC's complaint follows the merger of DoubleClick and Abacus Direct, the country's largest catalog database firm. DoubleClick has announced its intention to combine anonymous Internet profiles in the DoubleClick database with the personal information contained in the Abacus database.[97]

For extensive resources about legal aspects of cookies, see the EPIC *Cookies* Web page.[98]

Web beacons, also called *Web bugs*, *single-pixel Graphics Interface Formats* (GIFs), or *clear GIFs*, are extremely small graphics files, generally invisible without magnification, that are used by online advertisers to gather marketing data. Web beacons function by recognizing cookies. When a user visits a Website bearing a company's Web beacon, that company can now learn that this is the same user who visited that company's own site earlier. Or, if the cookie was placed by an advertising company, the company can identify each site that the user visits that bears the company's Web beacons. This allows company to know more about users' exact Internet usage and

surfing patterns than were possible using only cookies. Unlike the case with cookies, browsers do not readily detect the presence of Web beacons, nor do Website privacy policies typically disclose the use of Web beacons.[99]

Their use has provoked increasing consumer uneasiness; therefore, companies should proceed with care when electing to use or host Web beacons.[100]

69.3.3.10 Location Privacy. An explosion in the use of wireless devices, many of which are equipped with Global Positioning System (GPS) devices that allow for the identification or monitoring of the location of the device, has spurred corresponding public concern over how to protect the privacy of location information. Great marketing opportunities are envisioned, such as being able to extend special offers to consumers who are in the physical vicinity of a store or restaurant, but many people also have identified the potential for abuses.

A GPS-enabled receiver is the device that is commonly available through commercial retailers, and used by the general public to assist in navigation. The civilian GPS receivers deliver precise velocity and timing information, and very accurate location information. This device by itself does not transmit the data received from the satellite network to remote locations, nor is it typically capable of storing data regarding its long-term historical movements. However, most cell phones are now GPS-enabled and can transmit location data to the service provider or other third parties (such as Apps—see, e.g., Foursquare). Cell phones can also be used to track users even if they are not GPS-enabled. See Cell Phone Tracking Methods.

Because the technology that allows for the collection and use of this information is so new, industry best practices have not yet developed, and little legal guidance is available. It is expected, however, that location privacy will be the subject of intense policy and legislative debate over the next several years. Location-based advertising raises particular concerns when consumers are expected to pay the cost of incoming messages. There are also questions about the appropriate legal standards to access location information for law enforcement purposes. Because of consumer and media sensitivity to issues of location privacy, companies should carefully consider the legal and public relations aspects of collecting, using, or transferring such information.

A technology that was garnering a great deal of media attention as this chapter was going to press in July 2013 was *Google Glass*, a wearable, networked computer that includes a camera and projects information in the user's field of view. The device can not only record information from the user's environment, it can provide information about recognized people in the field of view.[101]

Great concern has been expressed by privacy advocates that information about individuals' every move will become available before any clear rules are established concerning the protection or use of that information. For example, EPIC writes,

> When individuals are moving about in public and private spaces, they do not expect to be tracked wherever they go. However, this expectation is being challenged as cell phones and other electronic devices now collect and store location data throughout the day. The expansion of location tracking technologies has significant implications for consumers and for their constitutional privacy rights.
>
> Over the last 10 years, law enforcement has stepped up its use of location tracking technologies, such as GPS (Global Positioning System) trackers and cell phones, to monitor the movements of individuals who may or may not be suspected of a crime. GPS is a geolocation network that consists of satellites and receivers that can calculate the precise location of a GPS device 24-hours a day (subject to environmental constraints). As of March 19, 2013, there are 31 satellites in the GPS constellation. The satellites and ground stations in the GPS network

are maintained by the U.S. Air Force Global Positioning Systems Wing. GPS satellites are designed to transmit three-dimensional location data (longitude, latitude, and altitude) as well as precise velocity and timing information to an unlimited number of users simultaneously. A GPS receiver is all that one needs to access the service. GPS satellites can not receive any data, they can only broadcast location and timing information.[102]

69.3.3.11 Genetic Discrimination. The collection and use of genetic information raises particular privacy concerns because this information can be used to determine whether individuals are prone to certain illnesses. This has implications for both employment and health insurance determinations. Among the most significant privacy bills recently enacted in the United States is the Genetic Information Nondiscrimination Act. The Act prohibits discrimination on the basis of genetic information with respect to health insurance and employment. From introduction to final passage, the bill took almost a dozen years.

69.3.3.12 Social Network Sites and Privacy. Among the most closely watched new areas of concern are the privacy impact of social network sites, such as Facebook and MySpace.[103] Concerns arise with respect to both the amount of personal information, particularly of minors, that is made available over the Internet, and also the advertising techniques that involve the collection and use of personal information. When Facebook introduced a revised News Feed service in 2006, the company was accused of invading the privacy of users and facilitating stalking. Over 700,000 users signed an online petition demanding the company discontinue the feature, stating that this compromised their privacy. Subsequently, Mark Zuckerberg, the CEO of Facebook, published an open letter apologizing for not having consulted with users prior to changing the service.

In 2007, the European Network and Information Security Agency issued a report on social networks and privacy. The paper concludes that social networks are a clear benefit to society; however, the study warns of the danger that new technologies pose in a world where there may be a false sense of intimacy created by social networks. The paper recommends government and corporate policy changes, increasing transparency of data handling practices, and encouraging social networking education rather than the banning of social networking sites in schools.[104] In 2008, the International Working Group on Data Protection in Telecommunications released a report and guidance on privacy in social networking services. The report identifies risks to privacy and security, and provides guidance to regulators, service operators, and users to counter these risks. Risks include the large amount of data collection; the misuse of profile data by third parties; insecure infrastructure and application programming interfaces. Regulators should ensure openness and require data breach notification. Providers must be transparent, live up to promises made to users, and use privacy friendly defaults. Privacy and consumer groups are also encouraged to raise the awareness of regulators, providers, and the general public.[105]

69.3.4 State Legislation. As noted in other sections, most federal privacy laws and regulations in the United States do not preempt state law. Most often, federal statutes set a "floor" on which states may create stricter rules. This approach to legislating is important because it reflects the federalism form of government that grants states significant power. It also allows the states to develop new approaches to emerging privacy challenges and to operate as "laboratories of democracy." Not surprisingly, in

the absence of substantial federal legislation, U.S. states have acted aggressively to pass new legislation regarding personal information.

The important point regarding state privacy laws is that such laws are numerous and often organized by industry or commercial sector. A company always should research the laws of the state(s) in which it operates to determine whether there are state-level laws that affect its use of personal information.

States already have passed legislation concerning emerging privacy issues, such as DNA databases, RFID tags, and other new technologies. Examples of significant privacy state laws follow.

- California passed an influential security breach notification in 2003. SB 1386 requires "any person or business that conducts business in California, and that owns or licenses computerized data that includes personal information, [to] disclose any breach of the security system . . . to any resident of California whose unencrypted personal information was, or is reasonably believed to have been, acquired by an unauthorized person." The statute imposes specific notification requirements on companies in such circumstances. The statute applies regardless of whether the computerized consumer records are maintained in or outside California. Many states subsequently passed security breach laws similar to California, and California subsequently extended the legislation in 2007 to cover medical records.
- California also enacted strong consumer privacy legislation in 2003. The California Financial Information Privacy Act provides the strongest financial privacy protections in the United States. Among the key provisions, the law gives consumers greater control over their personal information held by businesses, requires simple and understandable notices, and includes strong penalties.[106]
- According to the National Conference of State Legislatures, the majority of state legislatures have taken steps to safeguard genetic information beyond the protections provided for other types of health information. The state of New Mexico, for example, grants the individual the right to access genetic information, requires consent for a wide range of purposes, and establishes specific penalties for violations of genetic privacy.[107]
- Several states have passed laws to regulate the use of RFID. Washington State passed a law in 2008 that makes it a felony to intentionally scan another person's identification remotely without his or her knowledge and consent, for the purpose of fraud, identity theft, or some other illegal purpose.[108] A 2006 law in Wisconsin prevents the forcible implementation of RFID tags into humans.

69.4 COMPLIANCE MODELS. Responding to the challenges presented by rapidly developing privacy laws and the threats posed by hackers, requires the confrontation of complex new issues. A number of different paradigms are emerging. Each is interrelated, but for purposes of explanation, it may be best to consider each individually. Real-world solutions to privacy issues will, in almost every case, involve the use of multiple overlapping compliance models. This section illustrates, by providing a broad overview, the various models that are emerging as themes of privacy law compliance.

69.4.1 U.S. Legislation. The primary compliance model for privacy protection is the enactment of law that makes clear the obligations of organizations that collect and use personal information and the rights of individuals who provide

personal information. Privacy legislation is typically based on Fair Information Practices that are intended to promote the flow of personal data but to do so in a way that makes clear the purposes for which personal information is collected and will be used. Privacy legislation typically includes a variety of enforcement mechanisms, including a private right of action that allows an individual, or group of individuals, to file a lawsuit when a provision of the law has been violated; enforcement by the state of attorney generals, which can be particularly important when there is a pattern or practice of violation; as well as action by the Federal Trade Commission, which has the power to investigate unfair and deceptive trade practices.

69.4.2 U.S. Federal Trade Commission Section 5 Authority. The United States Federal Trade Commission has authority to investigate unfair or deceptive trade practices under Section 5 of the FTC Act. The FTC has applied this authority in a variety of privacy cases, involving primarily negligent security practices.

69.4.3 Self-Regulatory Regimes and Codes of Conduct. Self-regulatory regimes, including industry codes of conduct, are emerging as a means for achieving compliance with fair information practices among businesses. Self-regulation has two appealing aspects: (1) it minimizes the need for government to devote resources to enforcement, and (2) it allows the greatest degree of flexibility of business practices—important to those who stress that overregulation will stunt the growth of electronic commerce.

At the same time, self-regulatory approaches have come under increasing scrutiny from privacy officials, consumer organizations, and political leaders, particularly in the European Union. The EU national data protection commissioners (the Article 29 Committee) opposed the Safe Harbor proposal as inadequate. The Trans Atlantic Consumer Dialogue found that there was insufficient enforcement.[109] The European Parliament challenged the adequacy of the Safe Harbor Arrangement with the United States.[110]

Codes of conduct are explicitly contemplated by the EU Data Protection Directive as one means of compliance. The directive encourages EU member states to foster codes of conduct within self-policing industries as a means of ensuring compliance with data protection laws across entire business sectors.[111]

69.4.4 Contract Infrastructure. Several privacy laws contemplate the use of contracts to define how personal information will be collected, used, and/or transferred. Although the specific uses of contracts vary from law to law, an underlying trend is emerging: The rules governing the entire "life cycle" of personal information will be supplemented and maintained by a network of contracts.

Personal information, under most privacy laws, is collected from an individual subject to the terms of how that information will be used and to whom it may be transferred. In the absence of law, industry standards, or other guidelines, the rules for how personal information may be used are defined by the policies given to the data subject. Preferably, these policies will incorporate all of the elements of fair information practices. Once a company has collected personal information, it must process and use the data pursuant to the rules. Since personal information often passes from the collecting company to its affiliates, service providers, and others, additional contracts must be executed among and between all of these entities to ensure that the rules originally defined in the notice are followed.

Companies should begin to think about a "chain of contracts" that starts with fair information practices for the data subject—essentially a contract between the company and the individual. The terms of that contract are then enforced, by means of other contractual relationships, as personal information moves to different business entities, across national boundaries, and to third-party businesses.

69.4.5 Synthesis of Contracts, Technology, and Law. Contracts alone are not a complete solution to privacy issues. Although indemnity provisions and contractual warranties may, on paper, ensure that personal information has been properly collected, used, and transferred, the practical realities of the situation may be otherwise. Some points to consider:

- Policing such contracts may be beyond the means or the inclinations of many businesses.
- Most businesses are not likely to sue their trading partners to enforce these contracts, except in the most grievous cases.
- Most consumers do not perceive themselves as having enough at stake to bring actions on an individual level, but the possibility of class-action lawsuits may exist.
- The security of personal data, once compromised, is difficult to resecure—it is hard to unring the bell.
- Finally, given the massiveness of the data files typically involved, the speed with which data can be moved, and the multiplicity of jurisdictions into which it can be sent, governmental enforcement may face a larger challenge than governments will be willing to pay for.

Accordingly, technological controls on the collection, processing, access, storage, and transfer become a key element in the management of personal information.

Technology is the implementation of policy and business logic. For instance, when a customer makes a choice not to be included on a mailing list, it is technology that will segregate that person's information and ensure that the choice is honored. It is technology that will ensure that information relating to all the people who opted out of a proposed use of their data is not transferred to a third party for marketing uses. And it is, therefore, technology that ensures that the third party receiving personal information for marketing uses can use the information in those ways without incurring regulatory or public relations damage. Even simple policing techniques, such as the use of dummy names to detect unauthorized uses by onward transferees, can be administered more effectively using modern information technology. In this example, contracts and technology combine to add value. Contracts set terms of use for personal information that eliminate legal risk; technology ensures that the terms are followed. The end product is a valuable asset—usable personal information, the currency of the new economy.

It is vital that technological infrastructures embody both the specific legal rules that a company must observe with respect to personal information and the contracts created to manage that compliance process. Likewise, it is vital that the contracts themselves embody the legal requirements, the needs of the business, and the limitations of available technology. Business processes must adapt to the realities of the law, but must utilize technology and contractual relationships in order to maximize profits. To manage personal information assets effectively, technologists and businesspeople must

begin to think about law, while lawyers must become familiar with technology and business.

69.4.6 Getting Started: A Practical Checklist. An outline of some of the practical steps that must be taken for an organization to begin to implement effective personal information management practices follows. This checklist is not all-inclusive but is designed to illustrate how the core privacy principles and different compliance models begin to interact in the management process.

- Achieve buy-in, at the highest level of the organization, to the idea that personal information management must be part of an organization's critical infrastructure.
- Perform due diligence to identify *all* types of personal information collected and the routes by which the data travel in and out of the organization.
- Identify all of the uses to which the information is put during its life cycle through collection, processing, use, transfer, storage, and destruction.
- Identify each law affecting the collection, use, and transfer of personal information to which the company is subject.
- Create an institutional privacy policy that accurately considers both a commitment to abide by various legal requirements and the legitimate business activities of the organization.
- Create supporting materials that educate employees and instruct on policy implementation.
- Implement consistent data transfer agreements with all data-trading partners, vendors, service providers, and others with whom personal information is acquired or transferred.
- Build privacy management into the organization's strategic planning, providing sufficient resources for personnel, training, technology, and compliance auditing.
- Hold employees accountable for implementation and compliance with the privacy policy and contract requirements.
- Consider innovative approaches to privacy protection and business development that limit or eliminate the collection of personally identifiable information.
- Periodically audit compliance.

69.5 FURTHER READING

Agre, P. E., and M. Rotenberg. *Technology and Privacy: The New Landscape*. MIT Press, 1998.

Bennett, C. J., and C. D. Raab. *The Governance of Privacy: Policy Instruments in a Global Perspective*. MIT Press, 2006.

Bygrave, L. A. *Data Protection Law: Approaching Its Rationale, Logic and Limits*. Kluwer Law International, 2002.

Garfinkel, S., and B. Rosenberg. *RFID Applications, Security, and Privacy*. Addison-Wesley, 2006.

Harvey, I., A. Cavoukian, G. Tomko, D. Borrett, H. Kwan, and D. Hatzinakos, editors. *SmartData: Privacy Meets Evolutionary Robotics*. Springer, 2013.

Hoepman, J.-H., and I. Verbauwhede, editors. *Radio Frequency Identification: Security and Privacy Issues: 8th International Workshop, RFIDSec 2012, Nijmegen, The Netherlands, July 2–3, 2012*. Springer, 2013.

Nehf, J. P. *Open Book: The Failed Promise of Information Privacy in America*. CreateSpace Independent Publishing Platform, 2012.
Rosen, J. *The Unwanted Gaze: The Destruction of Privacy in America*. Vintage, 2000.
Rosen, J., and B. Wittes, editors. *Constitution 3.0: Freedom and Technological Change*. Brookings Institution Press, 2013.
Solove, D. J., M. Rotenberg, and P. J. Schwartz. *Information Privacy Law*. Aspen, 2006.
Solove, D. J. *Understanding Privacy*. Harvard University Press, 2010.
Solove, D. J. *Nothing to Hide: The False Tradeoff between Privacy and Security*. Yale University Press, 2013.
Stewart, D., editor. *Social Media and the Law: A Guidebook for Communication Students and Professionals*. Routledge, 2012.
Ziegler, K. S. *Human Rights and Private Law: Privacy as Autonomy*. Hart, 2007.

Websites

American Civil Liberties Union: www.aclu.org
Electronic Privacy Information Center: http://epic.org
European Commission, Justice and Home Affairs, Protection of Personal Data: http://ec.europa.eu/justice/data-protection/index_en.htm
Federal Trade Commission, Privacy Initiatives: http://business.ftc.gov/privacy-and-security
National Conference of State Legislatures, NCSL 50—State Legislative Tracking Resources: www.ncsl.org/legislative-staff.aspx?tabs=856,34,736
Privacy, Business and Law—Pandab is an online newsletter summarizing the top news articles on privacy, law and business: www.pandab.org
Privacy Exchange—a global information resource on consumers, commerce, and date protection worldwide: www.privacyexchange.org
Privacy International—Privacy International (PI) is a human rights group formed in 1990 as a watchdog on surveillance and privacy invasions by governments and corporations: www.privacyinternational.org

69.6 NOTES

1. In this context, various non-U.S. legal systems, especially European law, tend to use the term *data protection*, while the law in the United States tends to use the term *privacy*. This chapter uses the terms interchangeably, although readers should be aware that *privacy* often is considered to be a broader concept than *data protection*.
2. "What Is the NSA Domestic Spying Program?" *NSA Spying on Americans*, Electronic Frontier Foundation, July 31, 2013, https://www.eff.org/nsa-spying/faq#37
3. "EPIC Urges US Supreme Court to Suspend NSA Domestic Surveillance." EPIC Alert 20.14, July 12, 2013, http://epic.org/alert/epic_alert_20.14.html
4. I. Chandrasekhar, M. Wardrop, and A. Trotman, "Phone Hacking: Timeline of the Scandal: The Key Events in the Phone Hacking Scandal, which Has Led to the Closure of the *News of the World*, Halted Rupert Murdoch's BSkyB Takeover Bid, and Prompted the Arrest of Several Key Figures," *The Telegraph*, July 23, 2012, www.telegraph.co.uk/news/uknews/phone-hacking/8634176/Phone-hacking-timeline-of-a-scandal.html
5. Lord Justice Leveson, *An Inquiry into the Culture, Practices and Ethics of the Press: Executive Summary,* http://media.ft.com/cms/643fca7e-3a2c-11e2-a00d-00144feabdc0.pdf

6. Parliamentary Reports What Price Privacy? The unlawful trade in confidential information May 13, 2006; What Price Privacy Now? December 13, 2006
7. According to the Australian Government, "The Privacy Amendment Act includes a set of new, harmonised, privacy principles that will regulate the handling of personal information by both Australian government agencies and businesses. These new principles are called the *Australian Privacy Principles* (APPs). They will replace the existing Information Privacy Principles (IPPs) that currently apply to Australian government agencies and the National Privacy Principles (NPPs) that currently apply to businesses.

   "Under the changes, there will be 13 new APPs. A number of the APPs are significantly different from the existing principles, including APP 7 on the use and disclosure of personal information for direct marketing, and APP 8 on cross-border disclosure of personal information." See www.oaic.gov.au/privacy/privacy-act/privacy-law-reform
8. E.g., the Fourth Amendment to the U.S. Constitution.
9. The U.S. Freedom of Information Act typifies the former trend; the constitutions of Argentina and South Africa typify the latter.
10. See Draft Charter of Fundamental Rights of the European Union, Charter 4487/00, Brussels, September 28, 2000: Article 7 ("Respect for Private and Family Life") of the charter states: "Everyone has the right to respect for his or her private and family life, home and communications." Article 8 ("Protection of Personal Data") states: "1. Everyone has the right to the protection of personal data concerning him or her. 2. Such data must be processed fairly for specified purposes and on the basis of the consent of the person concerned or some other legitimate basis laid down by law. Everyone has the right of access to data that has been collected concerning him or her, and the right to have it rectified. 3. Compliance with these rules shall be subject to control by an independent authority."
11. Sweden: The Data Act of 1973, found in Svensk Författningssamling (SFS; Swedish Code of Statutes) 1973: 289, amended at SFS 1982: 446; United States: The Privacy Act of 1974, P.L. 93-579, 5 USC §552a; Germany: The Federal Data Protection Act (BDSG), available in English in Ulrich Dammann, Otto Mallmann, and Spiros Simitis, eds., *Data Protection Legislation: An International Documentation* (Frankfurt am Main: [West] Germany, 1977); France: Loi No. 78-17 du 6 janvier 1978 relative à l'informatique, aux fichiers et aux libertés, *Journal Officiel de la République Française*, January 7, 1978.
12. ETS No. 108, Strasbourg, 1981. Available at: www.coe.fr/eng/legaltxt/108e.htm (URL inactive).
13. Paris, 1981. An important distinction in the activities of the OECD (as opposed to the activities of the COE) is the involvement and support of the U.S. Government. The OECD is supported by nearly 30 governments. For more information, see www.oecd.org
14. Council of Europe, Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data (ETS No. 108): http://conventions.coe.int/Treaty/Commun/ChercheSig.asp?NT=108&CL=ENG
15. For example, Hong Kong, Estonia, and Lithuania seem to have based their privacy laws, in part, on the OECD Guidelines. Brazil, Malaysia, and Morocco are currently considering passage of privacy laws based on the OECD Guidelines. In April 2013, Uruguay ratified the Guidelines and they entered into effect on the 1st of August of that year.

16. See www.oecd.org/internet/ieconomy/oecdprivacystatementgenerator.htm
17. Directive 95/46/EC of the European Parliament and of the Council of October 24, 1995, on the protection of individuals with regard to the processing of personal information and on the free movement of such data, text available at http://europa.eu.int/eur-lex/en/lif/dat/1995/en_395L0046.html (URL inactive).
18. "Status of implementation of Directive 95/46 on the Protection of Individuals with regard to the Processing of Personal Data," European Commission, July 16, 2013, http://ec.europa.eu/justice/data-protection/law/status-implementation/index_en.htm
19. "Commission Proposes a Comprehensive Reform of the Data Protection Rules," European Commission Website, January 25, 2012, http://ec.europa.eu/justice/newsroom/data-protection/news/120125_en.htm
20. "Commission Proposes a Comprehensive Reform of Data Protection Rules To Increase Users' Control of Their Data and To Cut Costs for Businesses," EUROPA Press Releases, January 25, 2012, http://europa.eu/rapid/press-release_IP-12-46_en.htm?locale=en
21. Hong Kong, New Zealand, Chile, Argentina, and Canada have passed such laws; Japan, India, Brazil, and Australia, among others, have proposed new conforming legislation.
22. The Czech Republic, Hungary, Poland, Latvia, and Estonia have already passed national legislation that mimics the EU model. Romania, Turkey, and Russia, among others, have proposed such legislation.
23. Directive Concerning the Processing of Personal Data and the Protection of Privacy in the Telecommunications Sector (Directive 97/66/EC of the European Parliament and of the Council of December 15, 1997), available at www.ispo.cec.be/legal/en/dataprot/protection.html (URL inactive).
24. See David Banisar, *Privacy & Human Rights 2000*, EPIC and Privacy International, available at www.privacyinternational.org/survey/phr2000/threats.html (URL inactive).
25. European Commission, "Proposal for a Directive of the European Parliament and of the Council Concerning the Processing of Personal Data and the Protection of Privacy in the Electronic Communications Sector," available at http://europa.eu.institutionalize/comm/information_society/policy/framework/pdf/com2000385_en.pdf (URL inactive).
26. European Data Protection Supervisor. https://secure.edps.europa.eu/EDPSWEB/edps/EDPS
27. "EU Data Protection Supervisor Slams 'Smart Borders'," NewEurope Online, July 19, 2013, www.neurope.eu/article/eu-data-protection-supervisor-slams-smart-borders
28. https://secure.edps.europa.eu/EDPSWEB/webdav/site/mySite/shared/Documents/EDPS/MembersMission/Members/13-07-31_vacancies_EN.pdf
29. Charles Warren and Louis Brandeis, "The Right of Privacy," *Harvard Law Review* 4 (1890):193.
30. See, e.g., *Cox Broadcasting v. Cohn*, 420 US 469 (1975); *Time, Inc. v. Hill*, 385 US 374 (1967).
31. See Restatement (2d), Torts, § 652.
32. 5 USC § 552a. 5 USC § 552 incorporates the Freedom of Information Act (originally passed in 1966).

33. The Fair Credit Reporting Act of 1970 (15 USC § 1681 *et seq.*) was arguably the first sweeping privacy law.
34. P.L. 93-579, congressional finding and statement of purpose, paragraph (b).
35. 18 USC § 3121 *et seq*.
36. See *Smith v. Maryland*, 442 US 735, 742–45 (1979).
37. 18 USC § 3122(b)(2).
38. See 12 USC §§ 3401–22.
39. 18 USC § 2721 (1994).
40. See *Reno v. Condon*, 528 US 141 (2000).
41. "Earthlink Says It Refuses to Install FBI's Carnivore Surveillance Device," *Wall Street Journal*, July 14, 2000.
42. For more details on monitoring and surveillance, see Banisar, *Privacy & Human Rights*.
43. See "Interception Capabilities 2000" (Part 2/5: "The State of the Art in Communications Intelligence [COMINT] of the Automated Processing for Intelligence Purposes of Intercepted Broadband Multi-language Leased or Common Carrier Systems, and Its Applicability to COMINT Targeting and Selection, Including Speech Recognition"), European Parliament, PE 168. 184 Vol. 2/5.
44. 47 USC §1001 *et seq*.
45. See *United States Telecom Ass'n v. FCC*, 227 F.3d 450 (D.C. Cir. 2000).
46. Council of Europe, Convention on Cybercrime, CETS No. 185: http://conventions.coe.int/Treaty/Commun/QueVoulezVous.asp?NT=185&CM=8&DF=5/19/2008&CL=ENG
47. Council of Europe, Convention on Cybercrime, CETS No. 185, Status of May 19, 2008 http://conventions.coe.int/Treaty/Commun/ChercheSig.asp?NT=185&CM=8&DF=5/19/2008&CL=ENG
48. See, e.g., www.cdt.org/international/cybercrime/001211cdt.shtml (URL inactive).
49. "Foreign Intelligence Surveillance Court," Federal Judicial Center | History of the Federal Judiciary, www.fjc.gov/history/home.nsf/page/courts_special_fisc.html
50. "Beyond FISA," Electronic Frontier Foundation | Surveillance Self-Defense, 2013, https://ssd.eff.org/foreign/beyond-fisa
51. "The USA PATRIOT Act." Electronic Frontier Foundation (2005-08-03), http://w2.eff.org/patriot
52. "How PATRIOT Threatens Online Civil Liberties." Electronic Frontier Foundation (2013), http://w2.eff.org/patriot/why.php
53. Financial Services Modernization Act of 1999 (Pub.L. 106–102, 113 Stat. 1338). Signed into law on November 12, 1999; codified at 15 USC § 6801 et seq.
54. The Federal Deposit Insurance Corporation (FDIC), the Office of the Comptroller of the Currency (OCC), the Office of Thrift Supervision (OTS), and the Federal Reserve Board issued joint regulations. The Securities and Exchange Commission (SEC), the National Credit Union Administration (NCUA), the Commodity Futures Trading Commission (CFTC), and the Federal Trade Commission (FTC) have each issued separate sets of regulations.
55. "Who Must Comply with the Privacy Rule?" *Frequently Asked Questions for the Privacy Regulation*, Federal Trade Commission, Section A, Question 1, December 2001, www.ftc.gov/privacy/glbact/glb-faq.htm

56. See 47248 *Federal Register*, Vol. 61, No. 174, Friday, September 6, 1996.
57. See 12 CFR Ch. 11 § 211.5(d).
58. Individuals with whom a "financial institution" has an ongoing relationship.
59. All individuals who seek financial products or services for personal, family, or household purposes, whether or not those products or services are ever actually obtained.
60. A degree of variation has emerged concerning required information security standards under GLBA. The SEC and CFTC regulations set forth very general requirements; the OTS, OCC, FDIC and Federal Reserve Board jointly issued detailed information security guidelines (66 *Federal Register* 8615, February 1, 2001); the NCUA also issued guidelines paralleling those issued by the banking agencies (66 *Federal Register* 8152, January 29, 2001); the FTC issued a proposed regulation that represents a significant simplification of the banking agencies' guidelines (66 *Federal Register* 46112, August 7, 2001).
61. *Frequently Asked Questions for the Privacy Regulation*, Federal Trade Commission, December 2001, www.ftc.gov/privacy/glbact/glb-faq.htm
62. 15 USC § 6501 *et seq*.
63. Kelly/Warner Internet Law & Defamation Practice, "COPPA: Children's Online Privacy Protection Act Summary," 2012, www.aaronkellylaw.com/coppa-childrens-online-privacy-protection-act-summary/
64. 16 CFR § 312.9.
65. See www.ftc.gov/bcp/conline/pubs/buspubs/coppa.htm (URL inactive).
66. P.L. 104–191, August 21, 1996.
67. The final rule was published December 28, 2000, at 65 *Federal Register* 82,461 (2000).
68. Available at www.hhs.gov/ocr/privacy
69. See 45 CFR § 164.506
70. www.hhs.gov/ocr/office/index.html
71. *Patient Safety and Quality Improvement Act of 2005 Statute and Rule*. Health Information Privacy, US Department of Health & Human Services, www.hhs.gov/ocr/privacy/psa/regulation
72. More information about HIPAA's privacy regulations can be obtained at www.hhs.gov/ocr/hipaa or at hipaadvisory.com (URL inactive).
73. "*HIPAA—Frequently Asked Questions*." Health Information Privacy, U.S. Department of Health & Human Services. www.hhs.gov/ocr/privacy/hipaa/faq/index.html
74. 47 USC § 551 (1984).
75. 18 USC § 2710 (1988).
76. *Harris v. Blockbuster*, Electronic Privacy Information Center, February 2, 2010, http://epic.org/amicus/blockbuster
77. "January 2002: Study Finds Agents Have Sought Records From 220 Libraries," *Profile: USA Patriot Act,* History Commons Website, 2002, www.historycommons.org/entity.jsp?entity=patriot_act
78. "The USA PATRIOT Act." American Library Association (2013), www.ala.org/advocacy/advleg/federallegislation/theusapatriotact
79. See Section 69.2.2.1 for more details.

80. The full text of the Safe Harbor Principles is available at www.ita.doc.gov/td/ecom/menu.html (URL inactive).
81. Additional information on TRUSTe and BBBOnline can be obtained at www.truste.org and www.bbbonline.org. Note that neither organization offers certification of offline privacy practices and is therefore not available as a complete solution to any company that engages in any type of offline collection or use of the personal information of EU residents.
82. The European Commission has deferred a review of whether the privacy provisions of GLBA are "adequate" pending implementation and possible further legislative action in the United States.
83. Although the Safe Harbor principles allow for adequacy determinations for U.S. business sectors that are subject to particular sectoral privacy laws, no determination has yet been made of the adequacy of any particular law.
84. See Introduction and Annex to Safe Harbor Principles: 15 USC § 45. Note that banks, savings and loans, and entities such as stockyards and meat packers are excluded from the scope of the FTC Act and are therefore not eligible for Safe Harbor membership.
85. See draft Safe Harbor FAQ No. 11: Dispute Resolution and Enforcement, available at www.ita.doc.gov/td/ecom/RedlinedFAQ11Enforc300.htm. (URL inactive). Note that entities that are eligible for Safe Harbor membership because they are subject to the jurisdiction of the Department of Transportation may be subject to enforcement actions brought by the department.
86. See draft Safe Harbor FAQ No. 11: Dispute Resolution and Enforcement.
87. The "Safe Harbor" came into effect in early November 2000. The informal stay of enforcement by EU Data Protection Authorities against U.S. companies was due to end in June 2001. Some enforcement actions have been brought against U.S. companies already; for instance, American Airlines was forced to change its practices with respect to passenger information by the Swedish authority and Microsoft was fined $60,000 in Spain.
88. *U.S.–EU Safe Harbor Overview*, http://export.gov/safeharbor/eu/eg_main_018476.asp
89. During 2001, such keystroke monitoring was the subject of increasing controversy. In the first case testing the use of such technology for law enforcement purposes, a federal judge in New Jersey recently ordered the U.S. government to reveal how it had performed keystroke monitoring of an alleged criminal. See Opinion and Order, August 7, 2001, *United States v. Nicodemo Scarfo et al.*, available at http://www2.epic.org/crypto/scarfo/order_8_7_01.pdf (URL inactive).

    In another incident, judges of the Ninth U.S. Circuit Court, based in San Francisco, ordered their information technology personnel to remove keystroke-monitoring software that had been installed on their computers by order of the Judicial Conference of the United States. See Neil Lewis, "Rebels in Black Robes Recoil at Surveillance of Computers," *New York Times*, August 8, 2001.
90. See, e.g., *Smyth v. Pillsbury Co.*, 914 F. Supp. 97 (E.D. Pa. 1996).
91. 18 USC § 2510.
92. 18 USC § 2511(2)(d).
93. See *Watkins v. L.M. Berry & Company*, 704 F.2d 577 (11th Cir. 1983); *Ali v. Douglas Cable Communications*, 929 F. Supp. 1362 (D. Kan. 1996); *US v.*

Murdock, 63 F.3d 1391 (6th Cir. 1995); *Deal v. Spears*, 980 F.2d 1153 (8th Cir. 1992).

94. Fla. Stat. 934.03-934.09; Md. Cts. & Jud. Proc.Code Ann. Sec.10-402 (1980). The Maryland law was the basis for the criminal investigation into Linda Tripp's recording of her conversations with Monica Lewinsky regarding then-president Clinton.
95. See www.johndoes.org for more information. (URL inactive).
96. See *Dendrite International, Inc. v. John Does 1-14, et al.*, New Jersey Superior Court (Morris/Sussex Vicinage), November 28, 2000, docket no. MRSC-129-00.
97. "EPIC Files FTC Privacy Complaint Against DoubleClick," *Cookies*, Electronic Privacy Information Center, http://epic.org/privacy/internet/cookies
98. *Cookies*, Electronic Privacy Information Center, http://epic.org/privacy/internet/cookies/
99. Free software that allows users to identify Web beacons on websites is available at www.bugnosis.com (URL inactive).
100. *Online Tracking and Behavioral Profiling*, Electronic Privacy Information Center, http://epic.org/privacy/consumer/online_tracking_and_behavioral.html
101. "Google Glass May Compromise Privacy of User and Others," KTVU TV, May 16, 2013, www.ktvu.com/news/news/special-reports/google-glass-may-compromise-privacy-user-and-other/nXtHx/
102. *Locational Privacy*, Electronic Privacy Information Center, http://epic.org/privacy/location_privacy
103. *Social Networking Privacy*, Electronic Privacy Information Center, http://epic.org/privacy/socialnet/default.html
104. "Security Issues And Recommendations for Social Networks," ENISA, October 2007, www.enisa.europa.eu/doc/pdf/deliverables/enisa_pp_social_networks.pdf (URL inactive).
105. International Working Group on Data Protection in Telecommunications, "Report and Guidance on Privacy in Social Network Services, 'Rome Memorandum'," March 2008, www.datenschutz-berlin.de/attachments/461/WP_social_network_services.pdf?1208438491
106. Privacy Rights Clearinghouse, "California Financial Information Privacy Act: Senate Bill 1, Speier & Burton," www.privacyrights.org/ar/SB1Info.htm
107. National Conference of State Legislatures, "State Genetic Privacy Laws," updated January 2008, www.ncsl.org/programs/health/genetics/prt.htm
108. "First of Its Kind RFID Law Prohibits Malicious Spying on Consumers," PRWeb Website, www.prweb.com/releases/2008/03/prweb803114.htm
109. Submission of the TACD, March 30, 2000 at: www.ita.doc.gov/td/ecom/Comments400/TACDComments1.htm (URL inactive).
110. European Parliament resolution on the Draft Commission Decision on the adequacy of the protection provided by the Safe Harbor Privacy Principles and related *Frequently Asked Questions* issued by the US Department of Commerce (C5-0280/2000–2000/2144(COS)): http://epic.org/privacy/intl/EP_SH_resolution_0700.html
111. Directive 95/46/EC, Article 27. The International Commercial Exchange provides an example of such an industry initiative; see www.icx.org for more information.

CHAPTER 70

ANONYMITY AND IDENTITY IN CYBERSPACE

M. E. Kabay, Eric Salveggio, Robert Guess, and Russell D. Rosco

70.1 INTRODUCTION 70·2
- 70.1.1 Identity Theft 70·2
- 70.1.2 Politics 70·2
- 70.1.3 Harassment and Cyberbullying 70·2
- 70.1.4 Spam 70·5
- 70.1.5 Anonymous and WikiLeaks 70·6
- 70.1.6 Chapter Plan 70·6

70.2 DEFINITIONS 70·7
- 70.2.1 Cyberspace 70·7
- 70.2.2 The Real World 70·7
- 70.2.3 Identity in the Real World 70·7
- 70.2.4 Anonymity and Pseudonymity in the Real World 70·8

70.3 SOCIAL PSYCHOLOGY OF ANONYMITY 70·9
- 70.3.1 Deindividuation Theory 70·9
- 70.3.2 Identity in Cyberspace 70·14

70.4 BALANCING RIGHTS AND DUTIES 70·17
- 70.4.1 Benefits of Anonymity and Pseudonymity 70·17
- 70.4.2 Privacy in Cyberspace 70·18
- 70.4.3 Privacy and Freedom in Virtual Worlds 70·19
- 70.4.4 Defeating Dataveillance 70·19
- 70.4.5 Disadvantages of Anonymity and Pseudonymity 70·21

70.5 SYSTEMS ANALYSIS OF ANONYMITY 70·22

70.6 IMPLICATIONS AND DISCUSSION 70·23
- 70.6.1 Individuals, Families, and Schools 70·23
- 70.6.2 Ethical Principles 70·24
- 70.6.3 Corporations and Other Organizations 70·26
- 70.6.4 Internet Service Providers 70·26
- 70.6.5 A Free Market Model for Identity in Cyberspace 70·27
- 70.6.6 Governments 70·28

70.7 CONCLUDING REMARKS 70·29

70.8 SUMMARY 70·29

70.9 FURTHER READING 70·30

70.10 NOTES 70·30

70.1 INTRODUCTION. As electronic communications technology is now commonplace among increasingly international populations of computer users, one of the most hotly debated questions is how to maintain the benefits of free discourse while simultaneously restricting antisocial communications and behavior on the Internet. The proliferation of pseudonymous blogging and anonymous or pseudonymous commenting and tweeting complicates and potentially degrades the level of discourse online. International and intercultural differences in communications complicate the issue of anonymity and identity online; what is viewed as freedom in some parts of the world is perceived as license in other communities. Conversely, what are conceived by some as attempts to impose civility on international discourse are sometimes rejected as gross interference in freedom of speech by others.

At the heart of much of the debate over the advisability and possibility of imposing limits on behavior in cyberspace is the question of identity. Some of the most egregious abuse of cyberspace seems reasonably to be attributable in part to the ease of concealing identity; using no names or false names, malefactors can often escape almost all of the consequences of their actions. Sophisticated attacks include propaganda meant to sway public opinion, whether in the political field or in commerce; anonymous criminal groups can mask nefarious and manipulative intent with a friendly and inviting countenance.

70.1.1 Identity Theft. The FBI and other U.S. federal agencies consider identity theft a serious issue for law enforcement and national security: "Identity theft not only victimizes the individuals whose identity is stolen, but it can also pose a public safety and national security risk. … It is imperative that law enforcement know the true identity of the individuals they encounter while protecting our communities."[1]

70.1.2 Politics. Under totalitarian regimes, masking identity while using the instantaneous communications capabilities of the Internet has permitted world awareness of atrocities in such countries as Egypt, China, and Syria. Anonymization of an identity becomes key to survival in such situations. In an acknowledgement of the power of instantaneous communications, China has instituted a user-registration process to thwart anonymization. In the January 18, 2012, edition of *The New York Times*, in an article entitled "China Expands Program Requiring Real-Name Registration Online," author Michael Wines writes:

> China will expand nationwide a trial program that requires users of the country's wildly popular microblog services to disclose their identities to the government in order to post comments online, the government's top Internet regulator said Wednesday.
>
> The official, Wang Chen, said at a news conference that registration trials in five major eastern Chinese cities would continue until wrinkles were worked out. But he said that eventually all 250 million users of microblogs, called weibos here, would have to register, beginning first with new users.
>
> Mr. Wang indicated that under the program, users could continue to use nicknames online, even though they would still be required to register their true identities.[2]

70.1.3 Harassment and Cyberbullying. Corporations and individuals can suffer serious damage to their interests from abuse of anonymous communications. The availability of anonymous, highly critical postings to discussion groups about specific companies has resulted in problems in hiring new staff and depression of stock prices.[3]

One of the biggest issues facing law enforcement is how to deal with cyberthreats, especially anonymous threats. Many of the current laws, while similar in nature, are defused in court over technicalities by shrewd lawyers. It has been stated that new laws covering these new threats are sorely needed, but that the system is about five years behind the times.[4]

In an early case of cyberharassment, an innocent Florida businessman, Bruce Hovland, was harassed in 1997 by thousands of phone calls from angry strangers who complained about junk email that threatened to bill their credit cards for almost $200 in return for nonexistent pornographic videos they had never ordered and did not want. Hovland was the victim of a deliberate smear campaign, probably by a malefactor who had refused to pay rent at Hovland's marina and who had lost his boat in a seizure as a result. The anonymous malefactor spammed the 'Net in Hovland's name and suggested that people call his business number collect. Hovland guessed that he lost about two weeks of business because his phones were ringing off the hook. Hovland pointed out that his case was relatively minor; he imagined the mayhem if an emergency number were posted on the 'Net in such a fraud. The case illustrates the difficulty for victims in finding an agency willing to receive and follow up on complaints about such outrageous and dangerous attacks.[5]

Anonymous abuse of individuals can cause great personal distress; in 1997, for example, an Annapolis, Maryland, woman was mail-bombed after she warned other writers about extortionate fees from an author's agency; her name, phone number, and address were posted on alt.sex groups on the Usenet and resulted in floods of offensive phone calls.[6]

A woman in Atlanta was appalled when someone posted a photograph of an unknown anonymous woman with the victim's name and contact information; she received calls from men who told her the announcement claimed she was offering free sex. A victim of such anonymous harassment founded WHO@ (Working to Halt Online Abuse) to help victims fight this oppression.[7] The CyberAngels, an offshoot of the Guardian Angels vigilante group, claim to be willing and able to help victims of such harassment.[8]

In Florida, two students posted a picture of one of their high school classmates on a Website. The boy was pictured dancing with his prom date—but the girl's head was replaced by the picture of one of their male teachers. An electronic voice on the site announced that the teacher "must die." The student was profoundly disturbed by the intimations of homosexuality, as were the teacher and his colleagues. However, the state attorney's office reluctantly concluded that there is no valid state statute making it illegal for someone to publish libelous information anonymously on the 'Net. Although Florida does have a law forbidding "anonymous publication of material that holds a person up to ridicule or contempt," legal experts concluded that such a limitation on speech is unconstitutional.

In an article entitled "Web Plays Incendiary Role in Ohio High School Rape Case," from the January 7, 2013, edition of the *Los Angeles Times*, author Matt Pearce describes:

> The two high school football players accused of rape will get their day in court. The city of Steubenville, Ohio, however, will have to fight some of its battles online, where news of the case began.
>
> It's the story of a horrifying accusation met with small-town side-taking and blown up into a national scandal. Trent Mays and Ma'lik Richmond, both 16, are accused of raping a drunk and unconscious 16-year-old girl at a party on the night of Aug. 11 while other partygoers tweeted and Instagrammed about the attack.[9]

Pearce describes the ensuing cloud of information:

> The viral nature of the news was not limited to tweets on the night of the incident. The high-profile criminal case spotlighted by social media has put not only the defendants on trial but the culture surrounding them. Some town residents have compared the attention to a public lynching; others have said that without the attention, nothing would have been done.

Further exacerbating the situation,

> Hacker collective Anonymous recently posted video of youths who were said to be Steubenville High School students and alumni joking about the rape victim.

As a result of the varied and inaccurate information feeding angry public sentiment in the small town of Steubenville:

> The latest development in the case came Saturday as Steubenville Police Chief William McCafferty and City Manager Cathy Davison set up a blog dedicated solely to the case to "disseminate the most accurate information."

In this case, we have an alleged crime brought to light through real-time tweets and Instagrams; if the authors of these communications were witnessing a heinous crime such as this, why did no one take action to stop the crime, rather than to produce real-time tweets about it? Clearly, as a society, we have work to do with our children and ourselves.

Twenty years ago, a case like this would have played out in a court of law. Today, it plays out in a public forum with no rights afforded to any of those involved. The subsequent court case will be old news, with those involved already judged in unregulated public forums. The actual court results will hardly be noted except by the alleged suspects and victim, as the attention of tweeters, bloggers, and the public will have moved on to the next shiny and salacious topic.

In recent years, much has been written about cyberbullying of young people and tragedies that occurred as a result. However, in an article from the *Huffington Post* (www.huffingtonpost.com/larry-magid/cyberbullying_b_2162759.html) entitled "Researchers Dispel Myths About Cyber Bullying," technology journalist Larry Magid writes:

> A lot has been written about cyberbullying and I've seen some articles claiming that cyberbullying is more prevalent and more severe than in-person bullying. Some even refer to it as an "epidemic."
>
> But, in a presentation at the Family Online Safety Institute (FOSI) conference in Washington last week, a leading researcher on youth risk reported that the popular belief is actually a misconception. Compared to that bullying that takes place at school and other in-person venues, online bullying is both less prevalent and less distressing.

Magid brings us two key concepts to consider when thinking about media, cyberbullying, and anonymity; the first is from Michele Ybarra, President of the Center for Innovative Public Health, who states:

> . . . 17 percent of youth have reported online bullying compared to 39 percent who have experienced it "in person." Ten percent have been bullied by phone, according to Ybarra, while 14 percent have experienced bullying via text messaging.

Magid also brings forth another of Ybarra's key points, that is, the prevalent use of electronic communications among teens:

> For 12–17-year-olds, 95 percent are online, 77 percent have a cellphone, 23 percent have a smartphone, 63 percent text daily, and 76 percent use social media. Only 6 percent of teens use email on a daily basis.

Armed with that knowledge, Magid continues:

> Statistics from the U.S. Department of Health and Human Services also show that bullying—while still a significant problem—has not become more prevalent over the past few years. Bullying among high school students (Grades 9–12) remained flat (20 percent) between 2009 and 2011, with a slight decrease among boys and a slight increase for girls. The report found that in 2009, girls were twice as likely (22 percent vs. 11 percent) to be "electronically bullied" than boys.

Challenges remain. While most kids aren't bullied, a significant minority are, and the numbers are higher for lesbian, gay, bisexual, or transgender (LGBT) youth and those who are "perceived as different from their peers, such as being overweight or underweight, wearing glasses or different clothing, being new to a school, or being unable to afford what kids consider 'cool,' " according to the government Website StopBullying.gov.

70.1.4 Spam. Anonymity is fundamental to the abuse of the 'Net practiced by many spammers. Almost all spam contains forged headers in an attempt to escape retribution; in some cases the forgeries name real domains. One of the most significant early cases began in May 1997, when Craig Nowak, a college student, chose flowers.com at random as the fraudulent return address for his first attempt at junk email.[10] In so doing, he was merely following the suggestions of the unscrupulous purveyors of spam-distribution programs, who usually advise their naïve users to forge the headers of their junk email. Unfortunately for his victim, flowers.com is a legitimate business whose owner received 5,000 bounced messages and plenty of abuse for supposedly spamming the world. The enraged owner of flowers.com, Tracy LaQuey Parker, launched a lawsuit for damages and was supported by the Electronic Frontier Foundation (Austin chapter) and the Texas Internet Service Providers Association. In late September 1997, the plaintiffs won a temporary injunction against Nowak and his company, preventing him from further use of the appropriated domain name. In November 1997, Judge Suzanne Covington imposed a fine of over $18,000 on the defendants and added a particularly significant passage in her judgment that clearly enunciates the damages caused by forgery of return addresses:

> The Court additionally finds that the Plaintiffs have suffered, and will continue to suffer if [the defendant is] not enjoined, irreparable harm in the form of diminution in value of Plaintiffs' domain name; the possibility that Plaintiffs' reputation will be damaged forever by unauthorized use of a domain name associated with them in the controversial and hated practice of Internet spamming; and service disruptions. The potential harm to the Plaintiffs cannot be adequately valued in damages, and therefore the Plaintiffs have no adequate remedy at law. … The Court further finds that Plaintiffs … suffered actual damages from the unauthorized actions of the Defendants, including lost time, lost income, lost business opportunities, and lost use of their respective computer systems.[11]

In light of the seriousness of these abuses of inadequate identification in cyberspace, system managers and others concerned with the continued success of the Internet as

an effective communications medium should consider the reasons for abusive behavior of anonymous individuals. Is such abuse an aberration particular to cyberspace, or are there precedents in history and in other areas of life that can provide insights to shape public and corporate policy toward identification in cyberspace?

See Chapter 20 in this *Handbook* for more details of spam.

70.1.5 Anonymous and WikiLeaks. Vigilante Internet groups (a good example of which is the loosely organized group Anonymous) use the shield of anonymity coupled with technical acumen and a propensity to use the world stage as a means of forwarding a particular political agenda or world view with little risk of legal prosecution. Anonymous has no central directorate or power structure; anyone can claim to be affiliated with this collective.

Even Anonymous members can be pushed too far—when colleagues stray from the path. From an article in the *Financial Times* entitled "Anonymous Withdraws Support for WikiLeaks," by Tim Bradshaw, we learn:

> WikiLeaks founder Julian Assange has lost another high-profile supporter as Anonymous, the hacking collective, said the whistleblowing site had been "ruined by egos."
>
> Anonymous, a loose group of internet "hacktivists," has been one of WikiLeaks' closest allies during the past two years, launching cyberattacks against the site's opponents and, some claim, becoming a source of its material.

The source of discontent according to Bradshaw is:

> "We have been worried about the direction WikiLeaks is going for a while," Anonymous wrote in an online statement published overnight on Friday. "In the recent month the focus moved away from actual leaks and the fight for freedom of information further and further while it concentrated more and more on Julian Assange … But WikiLeaks is not—or should not be—about Julian Assange alone."
>
> The final straw for the shadowy group, which has also supported the Occupy Wall Street movement and the Arab Spring while at the same time attacking companies including Sony and News Corp, appears to have been Mr Assange's dinner this week with pop star Lady Gaga. …
>
> "That's great for him but not much of our interest," Anonymous said. "We are more interested in transparent governments and bringing out documents and information they want to hide from the public."

So if we can extrapolate from the above, it appears that even those who employ tactics of anonymity, be they groups or Julian Assange prior to his discovery, have lines that should not be crossed. Once the line is crossed, justice will be meted out.

70.1.6 Chapter Plan. Understanding anonymity online and its intended and unintended consequences is a topic that should be understood by all Internet users and consumers of communications that either originate on the Internet, or traverse the Internet.

In a corporate world, a CISO, CSO, and Human Resources departments should form a partnership to educate their constituency into becoming intelligent and savvy consumers of information, as their organizations will benefit greatly from critical thinking. The ability to understand, analyze, and discuss issues related to anonymity in an accountable society should be part of every citizen's tool kit.

This chapter reviews some of the findings of social psychology that show how and why anonymity has generally been associated with anti-social behavior. In the political

sphere, anonymity can be used to forward extremist agendas without accountability to fan the flames of hate or distrust based on falsehoods. It appears that anonymity on the 'Net will inevitably continue to spawn antisocial behavior in cyberspace and that we must somehow integrate this kind of abuse into plans for the further development of the Internet.

This chapter also touches on some positive uses of anonymity such as the use of anonymity in communications from inside repressive regimes to expose atrocities to the world. Additionally, some practical suggestions on how different degrees of tolerance for anonymity can be integrated into a cyberspace polity are discussed.

70.2 DEFINITIONS. To explore anonymity in cyberspace, it is helpful to establish some common vocabulary. What is cyberspace? What are identity and its absence? What is pseudonymity?

70.2.1 Cyberspace. In this chapter, *cyberspace* means the totality of electronic data storage and transmission; this chapter focuses on communications using the Internet. On the Internet, there are users of specific domains defined in the Domain Naming System (DNS), such as companies (those using addresses ending in *.com*); universities and other educational institutions (*.edu* addresses); U.S. Government agencies and departments (*.gov*); the U.S. military (*.mil*); and network service providers (*.net*); among others.[12] There are many geographical domain names, such as those ending in *.de* (Germany) or *.uk* (United Kingdom), which include users from commercial, educational, government, and military organizations. In addition, there are many Internet Service Providers (ISPs) in the *.net* and *.com* domains in the United States, and others throughout the world in geographical domains, whose members communicate through the Internet. Customers of value-added networks (VANs) such as America Online (AOL), MSN, Comcast, Cox Communications, Time Warner Cable, and a host of others communicate through the Internet as well as having restricted areas within their VANs where only members can post messages and files. Additionally, social media sites like Facebook, MySpace, LinkedIn, Google+, and Live.com can be manipulated to allow users of the service to post with real or imaginary names and as well to limit (or not) their audiences. All Internet users can post messages in public discussion lists on the Usenet or through remailing lists, which broadcast all inbound email to participants. Henceforth, for convenience, *the Internet* or *the 'Net* will include any of these users. Users of direct-dial bulletin board systems and modem-to-modem direct links are explicitly excluded from this discussion. For more details of data communications, see Chapter 5 in this *Handbook*.

70.2.2 The Real World. As used here, *the real world* (sometimes referred to as *meatspace*) refers to the material world of everyday human interactions. Using *real world* in this way is not intended to imply that cyberspace is less significant, less useful, or even less real; it is merely a convenient reference to distinguish the physical from the electronic.

70.2.3 Identity in the Real World. The key meanings of the noun *identity* for our purposes are:

- The collective aspect of the set of characteristics by which a thing is definitively recognizable or known

- The set of behavioral or personal characteristics by which an individual is recognizable as a member of a group
- The distinct personality of an individual regarded as a persisting entity; individuality

Roger Clarke summarizes the importance of personal identification in this way: The purposes of the interchange of identification include:

- to provide a gesture of goodwill,
- to develop mutual confidence, and
- to reduce the scope for dishonesty;
- to enable either person to initiate the next round of communications; and
- to enable either person to associate transactions and information with the other person.[13]

In line with these points, apparently ancient knights would tip their visor up for two reasons: for recognition as friend or foe, and out of respect for those being met.

70.2.4 Anonymity and Pseudonymity in the Real World. *Anonymity* can be defined simply as being without a name or with an unknown name. Someone who never uses the same unique identifier twice is anonymous. Notes left without any name are anonymous. Someone using a standard or canonical identifier ("student," "member," "guest," and so on) is also anonymous.

Pseudonymity is the *consistent* use of a false name that is not linked to personal details. When a musician chose to perform exclusively using the name *Sting*, it could have been an example of pseudonymity if the singer had concealed his usual name; however, in fact, Gordon Matthew Thomas Sumner never concealed his usual identifier.

These terms are imbued in English with negative connotations; nonetheless, anonymity has an honorable history in world philosophy and politics. In the United States, for example, the seminal *Federalist Papers*, which appeared in 1787 under the pen name "Publius," is held up as an outstanding example of anonymous contribution to political thought.[14]

Anonymity is protected in some government-sponsored monitoring programs, such as those supported by the Occupational Safety and Health Administration (OSHA) of the U.S. Department of Labor, where regulations stipulate for example that "Completed NETS Cards [hazard reports] can be given to a department manager or deposited, anonymously, into one of the drop boxes located throughout the facility."[15]

Clarke explores the concepts and history of human identity in a section of his paper on management and policy issues relating to human identification. Individuality has been a central concept in Western civilization since the Renaissance, says Clarke.[16] However, individuals can adopt more than one identity; for example, some women use their husband's surname in private life but maintain their original family name in their professions. Some people have several identities; for example, novelists with different styles sometimes use various pen names. The Danish philosopher Søren Kierkegaard wrote under 16 pseudonyms; Charles Dodgson wrote as Lewis Carroll; Eric Blair wrote as George Orwell. As the work of these writers illustrates, anonymity and pseudonymity are not inherently linked to anti-social behavior.

David G. Post wrote in a 1995 article,

> Should the government—can the government—control anonymity in cyberspace? Anonymity is pervasive in the "real world," to be sure; most of the encounters we have each day—on the bus, in the coffee shop, while shopping, etc.—are probably anonymous ones. But anonymity has perhaps never been quite so easy to achieve, and we have probably never been able to undertake so wide a range of anonymous transactions, as in cyberspace. For example, there are a number of "anonymous remailer" services now available on the Internet, computer systems where one can send an electronic mail message (addressed, say, to a friend across the country, to a list of individual addressees, or to an electronic bulletin board for public posting) in order to have the information identifying the computer from which the message was sent entirely stripped off and destroyed; the message, when it arrives at its destination, only contains information regarding the remailer from which it was sent, not the originator of the message. …
>
> And anonymity is not without its costs. Most obviously, a truly anonymous regime makes it impossible for "law enforcement"—broadly defined to include both public and private enforcement of legal rules—to obtain information about the persons responsible for harmful behavior so as to punish wrongdoers or compensate the victims of that harm. Indeed, the wider availability of anonymous message capabilities may actually increase the amount of harmful activity; if, as the economists might say, individuals need not "internalize" the costs of the harm imposed on others by their actions because they will not be called to account, we might expect more such behavior to emerge.
>
> But there are benefits—possibly profound benefits—as well. There are, of course, perfectly legitimate, indeed laudable, reasons why individuals might want to engage in communicative activity without putting any aspects of their identity at risk; think, for instance, of a discussion group among HIV-positive individuals, or battered spouses. And think, in the example most likely to prove persuasive when this question arises in the constitutional context, of the Federalist Papers, arguably the greatest single work of political theory in history that may well never have seen the light of day had its authors had to reveal their true identities.[17]

70.3 SOCIAL PSYCHOLOGY OF ANONYMITY. The popular mythology of the "Wild West" assumes that when large numbers of people migrated into western territories of the United States in the nineteenth century, they settled into new communities where few or no people knew about their neighbors' background. The anonymous or pseudonymous stranger is a staple of fiction; the assumption is that the lack of social ties increased the occurrence of lawlessness.

Technological change can have profound consequences on social behavior. For example, the development of mass-produced automobiles made possible the development of suburban shopping malls in the United States, which in turn led to an adolescent mall culture unimaginable in the 1920s.[18]

It seems quite likely that the pervasive spread of the Internet and the nature of anonymous and pseudonymous interactions is having equally profound effects on social organization and interactions.

The sections that follow review some well-established information on anonymity and social behavior from the social psychology literature. The implications of these principles for individuals, corporate policy makers, ISPs, and governments are discussed in the final section.

70.3.1 Deindividuation Theory. What do scientists know about the behavior of anonymous people? In general, the findings are not encouraging for the future of cyberspace unless we can somehow avoid the known association of antisocial behavior and anonymity. Early work on people in groups focused on anonymity as a root of the perceived frequency of antisocial behavior.[19] The anonymous members of a crowd show reduced inhibition of antisocial and reckless, impulsive behavior. They are subject

to increased irritability and suggestibility. One wonders if the well-known incidence of flaming (rude and largely *ad hominem* communications through email and postings on the Usenet and other public areas) may be traceable to the same factors that influence crowd behavior.

Later social psychologists formulated a theory of deindividuation[20] in which they proposed that one's personal sense of identity can be overwhelmed by the sense of belonging to a group.

Zimbardo suggested that anonymity, diffusion of responsibility, and arousal contributed to deindividuation and anti-sociality.[21] He noted that deindividuated people display reduced inhibitions, reduced reliance on internal standards that normally qualify their behavior, and little self-awareness.

70.3.1.1 Deindividuation and Technology. As mentioned briefly earlier, there is some reason to suppose that technology can contribute to the deindividuation of its users.

Anonymity has been postulated to account in part for the strong contrast between normal behavior and the behavior of those who become aggressive and hostile when driving cars. It seems intuitively plausible that being isolated in a tight personal space, a cocoon of glass and metal, gives some drivers a feeling of power precisely because of their (possibly temporary) anonymity. In addition, the anonymity of the other drivers may lead to a kind of dehumanization of the other. It would be interesting to study how many angry drivers refer to "that car" instead of "that driver" when they rail against some random act of road rudeness. Similarly, the isolation of an Internet user also may contribute to aggressivity; the object of wrath may, much like the driver of another car, be dehumanized.[22]

Sometimes it seems that email flamers are engaged in their version of a video game; they give the impression of losing sight of the real human beings on the other end of their verbal aggression.

Writers of computer viruses and others in the criminal computer underground may also focus so intensely on the challenge of defeating machines that they lose sight of their human victims. A commonplace informal interpretation of the insouciance of hackers and virus writers is that they are subject to the *video-game syndrome*: They seem to focus on their actions as if they were part of a game with only computers on the receiving end.

Criminal hackers have also expressed themselves as attacking systems, not people. At criminal-hacker conferences, it is commonplace to hear comments such as "Oh, I would never steal anything from a person, but if I found a radio in an office and it were labeled with a company sticker I wouldn't think twice about taking it."

Chapter 12 of this *Handbook* provides more detail on deindividuation theory.

70.3.1.2 Containment Theory. Walter Reckless[23,24] wrote about containment theory in a book entitled *The Crime Problem*. The theory addresses why people do not succumb to the pressure to commit crimes.[25,26] The theory includes external and internal containments that influence a person when making the decision to commit or not commit a crime.[27,28,29,30,31]

> External containments are external controls that "includes parental and school supervision and discipline, strong group cohesion, and a consistent moral front,"[32] "power of groups to make sure members comply with expected behavior,"[33] "external limits placed on children,"[34] or "forces that provide discipline and supervision including parents, police, schools, and the criminal justice system,"[35]

> Internal containments are a result of such things "as a positive self-concept, having a sense of direction in life, ability to tolerate frustration, and investment in and retention of group norms,"[36] "self-esteem, a strong sense of responsibility, internalized moral codes, tolerance of frustration, and positive goal orientations,"[37] or as simply "a strong conscience or a 'good self-concept,' "[38]

The inner containments are more important than external containments in determining whether a person commits a crime.[39] According to Taylor and colleagues,[40] containment theory consists of four processes: (1) "inner pushes and pulls lead an individual toward committing crime," (2) "inner containments inhibit criminal behavior … internal personal controls that lead someone to not commit criminal acts," (3) "outer pressures and pulls which lead to criminal behavior," and (4) "outer containments inhibit criminal behavior."

Deviant actions that are contemplated but not attempted in the physical realm may actually be more readily attempted in cyberspace. This assumption is exemplified by a person who would like to rob but is not right for the job; for example, a five-foot-tall 98-pound person who is afraid of guns would not attempt to rob a six-foot 200-pound bodybuilder carrying a handgun. But whatever motivates this 98-pound person could be executed via the digital world of computer crime, because the 98-pound person would not have to face his/her victim and the accompanying physical risk at the crime site. Another way to view this increased execution of the crime is that the criminal actor achieves physical distance, possibly thousands of miles from the victim, and perceives an anonymity enabled by the use of a computer. In reality, computers do not provide anonymity; computers and the network infrastructure keeps many details which can betray the criminal actor. It is the perception of anonymity that facilitates the actions.

70.3.1.3 General Crime Theory. General crime theory also provides some insight into the effect of anonymity and computer usage. General crime theory is an extension of Hirschi's social bonding with some elements of rational choice theory.[41,42] Social bonding theory suggests that social control is due to people avoiding crime because it will damage their relationship with family and friends.[43] Social bonds are what prevent people from committing crimes and consist of four elements: attachment, commitment, involvement, and belief.[44,45,46]

When people interact primarily through pseudonyms or completely anonymously, their attachments may be limited, their commitments may be limited, their involvement in social structures may be weakened, and their belief in immunity to the consequences of their actions may facilitate anti-social behavior.

70.3.1.4 Social Learning Theory. Social learning had its start when Burgess and Akers[47] combined Sutherland's differential association theory with operant and respondent conditioning.[48,49,50,51,52] Akers[53] modified what Burgess and he had done by replacing operant conditioning with the social learning of Bandura.[54,55,56] Akers[57] states that social learning theory also incorporates neutralization theory; the neutralization techniques are used to create definitions favorable to deviant behavior.

Sutherland's differential association theory consisted of nine propositions premised on the belief that deviant behavior action is learned from the individuals one associates with; not the inability to obtain economic success.[58,59,60,61,62,63] The nine propositions are based on three interrelated concepts of normative conflict, differential association, and differential social organization.[64] Differential association addresses normative or cultural conflict and how this conflict influences whether or not an individual commits a crime or deviant behavior.[65,66,67]

A crime/deviant action is committed when the definitions favorable to crime/deviant behavior outweigh the definitions unfavorable to crime/deviant behavior.[68,69,70,71,72,73] Matsueda[74] emphasizes that "not all definitions receive equal weight … rather, each is weighted by four modalities: frequency, duration, priority, and intensity." Definitions are dependent on larger societal context and are thus the definitions are subject to change over time.[75]

70.3.1.5 Anonymity and Aggression. Sometimes, anonymous people go beyond verbal abuse and seem willing to inflict harm on others. Experimental work by Zimbardo suggested that anonymity can significantly decrease barriers to aggression.[76] For example, when women were asked to deliver electric shocks to victims, those who agreed to wear white lab coats and hoods administered what they thought were longer shocks to the supposed victims compared with women who wore their own clothes and nametags.

In a cross-cultural study, Watson analyzed the correlations between the ritual, anonymizing costumes and war paint of warriors and their style of battle and their postbattle treatment of prisoners.[77] He found a strong positive relationship between anonymity and brutality.

Violent soccer fans seem to be disinhibited in part because of the anonymity of the crowd.[78] In an undercover investigation, a journalist found that anonymity and pseudonymity are integral components of the anti-social behavior of soccer hooligans.[79]

These findings suggest that so-called dark-side hackers may be influenced significantly in their willingness to cause damage to computer systems and networks precisely because their very anonymity influences them to cross normal behavioral boundaries. These people may not be the permanently, irremediably damaged human beings they sometimes seem; they may, instead, be relatively normal people responding in predictable ways to the absence of stable identification and identity.

70.3.1.6 Anonymity and Dishonesty. Does anonymity increase the likelihood that people will transgress rules and laws? Apparently, yes.

In an experiment involving children, young trick-or-treaters were asked to take only one candy from a bowl and then left alone or in groups, supposedly unobserved. Those children who had given their names to adults were observed to be far less likely to take extra candy or to steal coins than those who had remained anonymous, even when the adults were apparently away.[80] (If these effects of anonymity on youngsters are consistent and widespread, they may contribute to the problems of system administrators who are under siege by underage hackers.)

70.3.1.7 Deindividuation and Self-Awareness. Why does anonymity change people's normal inhibitions and influence them to behave abnormally? It seems that the deindividuation of anonymous people lowers their self-reflective propensities.

Exploration of the inner world of deindividuated people suggests that they are less aware of themselves and may even enter a state of altered consciousness.[81] Prentice-Dunn and Rogers studied the behavior of college men who, as in the work of Zimbardo cited earlier, were asked to administer what they thought were electric shocks to confederates of the experimenters who were masquerading as victims.

- Some subjects were subjected to dim lighting and loud background noise; their names were not used; and they were told that the levels of the shocks they gave would not be recorded. These subjects were thought to be deindividuated.

- Other subjects experienced bright lights in quiet rooms; they were called by name; and they were told that the shock levels they delivered would be monitored.
- The deindividuated subjects administered more severe shocks to their victims than did the individuated students.

These observations may tie into the work on autotelic experiences.[82] Autotelic experiences are deeply satisfying activities that result in a temporary loss of self-awareness; they typically occur in repetitive, challenging, feedback-rich activities such as programming (or perhaps virus writing) and criminal computer hacking. Csikszentmihalyi studied people in a variety of work environments and in their home life and hobbies. The subjects reported on their attainment of a state of timelessness, where the passage of time was insensible. Many who have programmed know how easy it is to forget to eat or to go home when deeply involved in work; similarly, writers and musicians can lose track of time. The research suggested that some of the key attributes of an activity that leads to this autotelic experience are rapidity of feedback (e.g., seeing an article grow as one writes or running an increasingly complex program under development) and being at the limits of one's abilities. In contrast, challenges that are too easy or too hard tend not to result in the loss of self-awareness that defines the autotelic state.

Several recent popular books dealing with criminal hackers have mentioned the ability of legendary hackers to stick to their hacking for hours on end as if they were entranced. Combine the autotelic nature of hacking with the deindividuation associated with anonymity, and we have a prescription for trouble.

70.3.1.8 Triple A Engine. Triple A engine refers to access, affordability, and anonymity of computing. Cooper used the Triple A Engine in examining online sexual activities (OSA) at the workplace; 20 percent of 40,000 adults he surveyed enaged in OSA (any activity involving sexuality). Cooper wrote that there are three different profiles of people who engage in OSA (bullets added):

- The first and most common is the Recreational User, who engages in OSA in a nonproblematic manner.
- The second type, the Sexually Compulsive User, refers to those individuals who have an established pattern of unconventional and maladaptive sexual practices. When faced with the Triple A Engine, these individuals often transfer the locus of their pre-existing problems to Internet-based activities. These online activities may in turn reinforce and exacerbate the problematic behaviors.
- Finally, the third profile was labeled the At-Risk User. These individuals have no prior history of sexual compulsivity, but they too often find their better judgment, coping strategies, and impulse control are overwhelmed by the power of the Triple A Engine. Research has revealed that some employees continue to access forbidden sites and sexually explicit material, in defiance of workplace policies.[83]

70.3.1.9 Anonymity and Prosocial Behavior. The picture is not necessarily all bad, however. Sometimes a different environment actually can liberate anonymous subjects from their counterproductive inhibitions. For example, in some cases, it can be shown that anonymity has an unusual effect: It increases prosocial behavior instead of increasing anti-social behavior.[84] Gergen, Gergen, and Barton put people into brightly lit chambers or in totally dark chambers and monitored the behavior of

the strangers they had put together; in the dark room, there was much more uninhibited and positive expression of physical contact such as hugs and of emotional openness such as discussions of personal matters.[85] However, directionality of these effects may be affected by the demand characteristics of the situation. That is, the way the experimental objectives were described could cause significant differences in the subjects' behavior.

The constructive, supportive communications often seen in discussion groups dealing with substance abuse, abusive relationships, and other personal and interpersonal problems illustrate the possible benefits of anonymity in a positive context.

70.3.2 Identity in Cyberspace. What exactly is meant by *identity* when using electronic communications? Is one's email address an identity? Could a made-up name be an identity?

Identity on the Internet is primarily the email address.[86] The email address sometimes provides crude and unreliable information about affiliation (e.g., domain names *.gov, .mil, .edu*) and geographic location (e.g., *.ca, .uk, .fr*).[87] Roger Clarke, a scholar with a long professional interest in questions of identity, identification, and privacy in cyberspace, has written an excellent introduction to these questions.[88] For discussions of information technology, Clarke defines *identification* as "the association of data with a particular human being."[89]

Identity now must also include Facebook, LinkedIn, MySpace, Google+, Live.com, sites such as Match.com, and any other host of Internet sites where identities real or imagined can exist. There are some identities in cyberspace where falsification and deception are actually part of the titillation of the business. Such a site is Ashley-Madison.com, where married individuals can register an identity of choice in order to arrange for an affair with another married person who too shares a desire for an extramarital affair.

Identities in other matchmaking sites, such as Match.com and eHarmony, are prone to falsification. In an article ("Online Dating Scams: Buyers Beware") from the *Huffington Post*, originally posted on February 14, 2012, and updated on May 15, 2012, author Ann Brenoffann writes:

> Online dating services are booming businesses, and the boomers who are joining them are doing so in big numbers. People aged 50 and older represent 25 percent of membership on the popular dating site Match.com—a 45 percent jump in the last five years, a spokesman said.
>
> But not everyone who hopes to find a mate online is falling blissfully in love. There have been a rash of complaints against online dating sites, according to the Better Business Bureau. And there is a class action lawsuit pending in the U.S. District Court that claims more than 60 percent of the profiles on Match.com are fraudulent—something Match.com spokesman Matthew Traub calls an "unfounded allegation in a money-seeking litigation."

Brennoffann continues:

> Then there is another, perhaps darker, side to the consumer complaints: People reporting that they were bilked out of money by those they connected with through an online dating site. The Better Business Bureau listed online dating sites among its top 10 scams for 2011, and one consumer advocacy group says its members have lost millions from online dating scams. … There is no attempt to verify the information someone posts—something the sites are generally upfront about.

Buyers beware indeed.

***70.3.2.1 Theory of* Nymity.** There has been considerable discussion on the 'Net about the kinds of identity (sometimes called *nymity* in these discussions) that people assume in cyberspace. One of the best-known writers on this subject was "L. Detweiler," apparently a pseudonym that has never been traced to any other identifier. Detweiler suggested that identity on the Internet is amorphous and unstable because there is no one-to-one relationship between a person and an email address. One person may use multiple email addresses, and many people may share a single address.[90]

Detweiler conducted a vigorous battle against what he perceived as a sinister and deceptive practice he called *pseudospoofing*. Pseudospoofing, in Detweiler's conception, is the use of multiple aliases by an individual or a conspiracy; these aliases allow the perpetrators to deceive observers into misjudging the number of people agreeing or disagreeing over specific positions.[91] Current examples of such a system are online travel, services, and restaurant sites. Sites such as Trip Advisor, Yelp, Chowhound, and the like must be vigilant that owners of questionable businesses do not flood the review site with favorable reviews using a multitude of false personas. Such abusive pseudonyms are currently known as *sock puppets* and can be used as elements of dishonest *online reputation management*. Other abuses of pseudonymity include:

> Astroturfing (named after the synthetic grass, "AstroTurf") is the act of creating fake grassroots organizations to endorse a client or a client's point of view or criticize a political position or a competitor under the guise of pseudo-authenticity. Typically, these organizations are staffed with people directly involved with the company or industry they're promoting. In a notorious case of astroturfing, a public relations firm was hired by a coal company to lobby against a bill in the U.S. Congress that would have regulated coal-industry activities. They forged letters—including forged logos—claiming to be from various social-activism groups and were eventually discovered.
>
> Sock puppetry is the practice of pretending to be someone else and using the same techniques as astroturfing, the difference being the scale (sock puppetry is mainly done with individual reviews). For a hilarious send-up of sock puppetry, see the article in "Uncyclopedia, the content-free encyclopedia." Just don't take any of the content seriously.
>
> Flogging is, in current marketing parlance, a portmanteau term made up from the two words "fake" and "blogging," and refers to the practice of companies employing fake bloggers to write glowing reviews of certain products.[92]

A current illustration of the problems of pseudospoofing is the widespread difficulty experienced in online voting. The ease with which identity can be created, coupled with the ease of automatically scripting multiple votes, leads to unreliable tallies in almost all online polls. The only way to avoid such abuses is to enforce some tight coupling of real-world identity with the electronic identity registered for voting. Eventually, biometric identification may be the only acceptable form of authentication for voting online.[93]

70.3.2.2 Types of Anonymity and Pseudonymity. To understand the problem of anonymity and pseudonymity, it is useful to define varying degrees of the behavior. Froomkin distinguishes among four forms of imprecise or absent identification:

1. **Traceable anonymity.** Any anonymous remailer that keeps a record of the relation between the original message and the anonymized message allows tracing.
2. **Untraceable anonymity.** No record is kept or available showing the origin of an anonymized message.

3. **Untraceable pseudonymity.** A continuous identity or persona allows communication with a correspondent, but there is no way to link the pseudonym to the correspondent's real-world identity.
4. **Traceable pseudonymity.** Someone, somewhere has the information required to complete the link between a pseudonym and a real-world identity.[94]

The anonymizing remailer *anon.penet.fi* was actually a *traceable pseudonym remailer*. When a message was received by the server, its headers were stripped and it was assigned a fixed random pseudonym. In order to allow replies to the pseudonym to be forwarded to the original sender, every pseudonymous identity was linked in a table to the original email address. When Finnish police ordered Johan Helsingius in 1995 to identify the pseudonymous poster of copyrighted Scientology texts on the Usenet, Helsingius felt obliged to reveal the link.[95] Traceable online anonymity allows people to maintain their privacy by using screen identities, but many ISPs will furnish the real-world identity to law enforcement officials with a warrant or to tort lawyers with a subpoena. AOL, for example, furnished subscriber details to the lawyers for a Caribbean resort considering a lawsuit for defamation based on postings by "Jenny TRR."[96] All ISPs that charge money for access inherently provide traceable pseudonymity even when they permit false screen names.

Larry Lessig, David Post, and Eugene Volokh also distinguish between anonymity, pseudonymity, and traceability in their Internet course on cyberspace law.[97] They emphasize that private organizations, such as Internet service providers, can freely set terms of service that allow or forbid anonymity, but they may be required in the United States to provide traceability if the courts require it for individuals.[98] In general, it is difficult to see how the rule of law can apply to cyberspace without some form of traceability. Whether civil law or the criminal law is involved, the defendant must be found for court proceedings to have any effect. Untraceable anonymity and untraceable pseudonymity preclude justice for the aggrieved parties. Clarke points out that privacy interests are always to be balanced with other interests, such as the public good, commercial interests, and the interests of other individuals.[99] For example, the desire to post anonymous libel (construed as privacy of personal behavior) conflicts with the desire of the victim to be free from libel; anonymity makes it impossible to use the civil law for redress.

70.3.2.3 Why Anonymity and Pseudonymity Are Commonplace in Cyberspace. Surely anonymity and pseudonymity are possible using older means of communication: People have thrown rocks through windows, sent anonymous vituperation through postal mail, and harassed people with anonymous phone calls for millennia, centuries, and decades, respectively. Historically, such behavior has been of relatively minor importance. How is it that anonymity and pseudonymity seem so important in cyberspace?

One factor is the ease with which one can be untraceably anonymous in cyberspace. Many ISPs allow users to define screen names or aliases; some ISPs, such as AOL, distributed millions of trial subscription disks that allowed one to create an identity and then dispose of it after 50 hours online. Spammers often use disposable free accounts from such services as Hotmail, even defeating methods designed to stop them from abusing the services.[100]

Some moderated Usenet groups and email-based mailing lists require a real name, but there is little effort or even possibility, as things currently stand, for authentication of

such names. Unmoderated groups, by definition, do not require real names at all; such groups allow postings by anyone. Almost all contributions to countercultural or frankly criminal Usenet groups and distribution lists are anonymous or at best pseudonymous.

In addition, some people systematically forge the headers of their email, often introducing completely false addresses for their origin but occasionally picking real addresses, whether by accident or design, as discussed earlier in the flowers.com case (Section 70.1.4).[101] Forging email headers can conceal the true origin of a message.[102] Junk email almost always includes false return addresses. Another way to generate false header information is to compromise someone's email account. If security on an email address is compromised, messages can be forged with a false ID. From a mobile media perspective, readily available software and services that allow phone and text spammers to alter the sending phone number displayed on mobile phones to one of a reputable or nonspam-associated telephone number are alive and well today—and harnessing the power of the cloud.

Anonymous remailers permit users to send email to a central address where all identifying information is stripped, and the original text is then retransmitted to the desired destination. In the wake of the tragic events of September 11, 2001, when terrorists killed thousands of people in attacks on the World Trade Center in New York City and on the Pentagon in Washington, D.C., commentators expressed doubts about the future of anonymizing services. "Using anonymizers at all raises all sorts of red flags," said John Young, operator of the Cryptome Website (www.cryptome.org), which covers intelligence matters. Even before the terrorist attacks, the well-respected Freedom Network operated by Zero-Knowledge Systems had experienced financial difficulties; that service shut down on October 22, 2001.[103]

In 2007, Dan Egerstad posted the user names and passwords for 100 email accounts of foreign embassies and human rights groups that used The Onion Router (Tor) network:

> Under Tor's architecture, administrators at the entry point can identify the user's IP address, but can't read the content of the user's correspondence or know its final destination. Each node in the network thereafter only knows the node from which it received the traffic, and it peels off a layer of encryption to reveal the next node to which it must forward the connection. (Tor stands for "The Onion Router.")
>
> But Tor has a known weakness: The last node through which traffic passes in the network has to decrypt the communication before delivering it to its final destination. Someone operating that node can see the communication passing through this server.[104]

The other factor that makes anonymity and pseudonymity especially significant in cyberspace is the ease of replication of messages. An anonymous communicator could hope to reach at best a few dozen or hundred people in a day with the phone system or by mail; in most areas, each phone call or letter would cost something to send. In contrast, spammers using botnets and those with large following on sites such as Twitter can reach countless people in seconds with very little effort and no cost at all.

70.4 BALANCING RIGHTS AND DUTIES. Is there a basis for evaluating the ethics of using anonymous and pseudonymous communications? Are these modes of communications protected by principles of privacy, for example?

70.4.1 Benefits of Anonymity and Pseudonymity. Historically, anonymity has been recognized as a benefit for acts of charity for those wishing to

remain nameless for political, social, or personal reasons. Maimonides, the great twelfth-century rabbi, defined anonymous giving to anonymous recipients as the second-highest level of charity.[105]

In discussions of whether society ought to restrict anonymity and pseudonymity, a common argument is that these modes of communication are necessary to fight tyrannical corporate and political institutions. Anonymity and pseudonymity are, in this view, expressions of the right to privacy. Electronic Frontier Foundation has filed First Amendment friend-of-the-court briefs in support of anonymous speech in which even the plaintiff later stated that "anonymous speech is worth protecting."[106] Abuses are the price society has to pay to preserve the benefits of these tools of expression. The next sections examine the concepts of privacy and the resulting benefits of anonymity and pseudonymity.

70.4.2 Privacy in Cyberspace. If privacy rights are claimed to protect anonymous and pseudonymous communications, it is important to understand the concepts of privacy.

Clarke defines privacy as "the interest that individuals have in sustaining a 'personal space,' free from interference by other people and organizations." He analyses the concept further, naming four dimensions of privacy:

1. **Privacy of the person.** Freedom from compulsory tampering with one's body
2. **Privacy of personal behavior.** Freedom in such matters as sexual preferences, religion, and politics
3. **Privacy of personal communications.** Freedom from routine monitoring of interpersonal communications
4. **Privacy of personal data.** Control over who can obtain and what can be done with personal information[107]

Political discussion groups, resistance to totalitarian regimes, and discussions of socially embarrassing or traumatic problems are made easier for many people by the use of pseudonyms or of anonymity. Anonymity permits unrestricted political speech, whistleblowing with reduced likelihood of retaliation, and public or private discussions of potentially embarrassing personal problems.[108] Ubois writes:

> Anonymous communications are helpful in many ways. They've long been a tool of suicide prevention hotlines, suggestion boxes, and personal ads. Anonymity assures privacy, confidentiality, and security for individuals, but it also highlights the clash of interests between the individual and the community. Under a repressive government, it is a vital tool for keeping discourse alive. Just consider that Tom Paine would have landed in prison shortly after the publication of Common Sense if his identity hadn't been kept a secret.[109]

In chat rooms and multiuser dungeons, anonymity permits a flowering of imaginative departures from the strictures of a participant's real-world identity, social status, personality, gender and gender preferences, political affiliation, national origin, and religion. Multimedia environments such as WorldsAway (http://secondlife.com/) provide an imaginative pseudonymity by allowing players to select a name and a pictorial representation of themselves (an avatar) with amusing and fanciful features, such as various imaginary animal heads, skin colors, body shapes, and so on. Players adopt personae that can be quite different from their real-world identities, yet there is a consistent identity within the virtual world. Because of this consistency, social mechanisms

have arisen in these worlds; for example, avatars can be invited to join parties if they are perceived as friendly or excluded and shunned if they have violated the norms of the imaginary world.

Anonymous, invisible electronic personalities can escape some of the damaging effects of intolerance and prejudice.[110] Everyone probably knows of Peter Steiner's famous *New Yorker* cartoon showing two dogs at a terminal, one of whom is saying "On the Internet nobody knows you're a dog."[111] Another current example is the country music video "Online" by Brad Paisley in which he sings "I'm much cooler online," and boasts about completely illusory qualities.[112]

For example, some professors may spend more time and effort in discussions with undergraduate students if they do not realize with whom they are corresponding.[113] At an intellectual level, stripping the authors of published materials of all details of their age, sex, race, national origin, and other attributes can reduce the effects of prejudice and focus discussion on substance. Absent such details, correspondents must perforce focus on the texts rather than on personalities.[114]

In electronic commerce, anonymity is a prerequisite for successful implementation of some trading systems.[115] All electronic or digital cash schemes (e-cash) emphasize the value of anonymous transactions for safeguarding consumer privacy.

In a legal sense, anonymity and pseudonymity are analogous to limited liability—a concept familiar from the business world. In business, people pool their assets into a limited liability partnership or other form of collectivity to prevent seizure of all their private assets if their collective entity becomes liable for debts or penalties.[116] Perhaps cyberspace anonymity and pseudonymity can encourage collective publications in an analogous way; for example, Post and others have formed the Cyberspace Law Institute. Some members would be reluctant to participate if their real-world identities were known to the public.

Thus, anonymity and pseudonymity cannot reasonably be forbidden without the loss of important benefits to individuals, corporations, and society at large.

70.4.3 Privacy and Freedom in Virtual Worlds. With the increasing growth of global access to the Internet, millions of people have joined virtual communities or *virtual worlds* where they can interact anonymously or pseudonymously. The freedom from discrimination based on physical, cultural, gender, and sociological attributes has proven liberating for many members of these virtual worlds; many participants have formed strong friendships that may even extend into the real world.[117] On the other hand, some critics argue that the avatars supplied by the owners of these virtual worlds reinforce conventional gender stereotypes.[118]

70.4.4 Defeating Dataveillance. Another area where anonymity and pseudonymity have marked benefits is in preventing intrusive monitoring of individual behavior in cyberspace. Clarke has defined dataveillance as surveillance using electronically compiled information about a person.[119]

The growth of some kinds of electronic commerce has increased pressures for strong identification and authentication[120]; anonymity serves to protect privacy in a world of electronic commerce. For example, without anonymous digital cash, it would be easy to accumulate detailed records of every electronic purchase made by an individual. Complete knowledge of purchasers' interests can be unfair for customers; for example, knowing that a user is addicted to fantasy simulation games, a retailer may neglect to offer that person a discount—or may even increase the price of the next game.[121]

Froomkin summarizes the issues well in his magisterial review of the challenges of anonymity in cyberspace:

> Anonymity lies at the heart of three interrelated problems arising from computer-aided communications over distributed networks (which I will call "the Internet" for short). First, communicative anonymity is an issue in itself: the Internet makes anonymous communication easy, and this has both good and bad consequences. …
>
> Second, the availability of anonymous electronic communication directly affects the ability of governments to regulate electronic transactions over the Internet (both licit and illicit).
>
> Third, anonymity may be the primary tool available to citizens to combat the compilation and analysis of personal profile data, although data protection laws also may have some effect. The existence of profiling databases, whether in corporate or public hands, may severely constrict the economic and possibly even the political freedoms of the persons profiled; although profiling may not necessarily change the amount of actual data in existence about a person, organizing the data into easily searchable form reduces her effective privacy by permitting "data mining" and correlations that were previously impossible.[122]

Froomkin discusses digital cash as an application of electronic anonymity and emphasizes the potential for abuse by "the Argus State" if anonymity is not guaranteed for readers. For example, he points out, in the absence of anonymous digital cash, reading texts on the Internet using micropayments for each access could provide a traceable record of a person's interests. Such records would be a gold mine for repressive regimes worldwide.

In June 2013, as this chapter was going to press, interest and debate were generated by the disclosure that the U.S. National Security Agency (NSA) had been using massive data capture under secret warrants issued by the Foreign Intelligence Surveillance Court to gather metadata about phone calls by hundreds of millions of American citizens. Decrying the abuse of the USA PATRIOT Act[123] to carry out dataveillance of such activities, the American Civil Liberties Union (ACLU) filed a lawsuit shortly after the revelations were published.[124]

These revelations suggested that anonymizing services that conceal all Internet activity, such as The Onion Router (Tor), would see an upsurge in use. The Tor Project describes its functions as follows:

> Tor is free software and an open network that helps you defend against a form of network surveillance that threatens personal freedom and privacy, confidential business activities and relationships, and state security known as traffic analysis.
>
> Tor protects you by bouncing your communications around a distributed network of relays run by volunteers all around the world: it prevents somebody watching your Internet connection from learning what sites you visit, and it prevents the sites you visit from learning your physical location. Tor works with many of your existing applications, including web browsers, instant messaging clients, remote login, and other applications based on the TCP protocol.[125]

Again, trying to ban anonymity and pseudonymity would have serious disadvantages for everyone, not just the benefits of impeding abuse by a minority of anti-social users.

Finally, police forces and intelligence agencies always recognize anonymity and pseudonymity as potentially important tools in specific investigations. Undercover or covert operations and agents can provide critically important evidence in anti-crime and anti-terrorist work. Even organizations such as the FBI-sponsored InfraGard recognize the value of anonymized information-sharing about attacks on computer systems.[126]

70.4.5 Disadvantages of Anonymity and Pseudonymity. Several commentators have reviewed the abuses of anonymous and pseudonymous modes of communication.

In considering the benefits of having a professor communicate with an anonymous student, misrepresentation of identity by such an undergraduate in one sense manipulates a professor into a decision based on a falsehood.[127] Contrary to the beliefs of many supporters of anonymity, social interactions are not necessarily equivalent to isolated streams of data interchange; conversing with a student can be enriched by having a sense of previous conversations, a picture of shared knowledge based on a relationship. Anonymity strips interactions of the deeper communication that is enhanced by such relationships.

Widespread use of untraceable e-cash may lead to increased fraud, tax evasion, money laundering, extortion, blackmail, and kidnapping.[128] Some crimes where solicitation leads to potential blackmail—for example, hiring a murderer—may become easier with anonymity and untraceable e-cash. Industrial sabotage by anonymous publication of trade secrets can damage organizations; for example, the publication of RC4 encryption algorithms from the respected company RSADSI has lowered the monetary value of the algorithm. Froomkin writes, "[The] inability to redress legitimate claims is, I believe, the strongest moral objection to the increase in anonymous interaction."[129]

Jurisprudence in the United States has generally supported claims to a right of anonymity in political speech. However, there have been several precedents where anonymity used to cloak socially harmful acts has been stripped from the perpetrators. Perfect (untraceable) anonymity prevents society from bringing sanctions to bear on malefactors.[130] Detweiler suggested "that Internet anonymity is a bad thing [and] that all user accounts should lead back to real users in the hopes of improving online behavior, especially in chat systems and the like." He responded to critics who claimed that "anonymity is an important part of life and ought to be part of the Internet as well" by pointing out that in real life, anonymous people cannot engage in such activities as opening a bank account, getting a driver's license, getting telephone service, or buying insurance coverage. He concluded, "So grow up and accept responsibility for what you do on the Internet."[131]

L. J. Rose, well known in cyberspace-law circles, writes scathingly of the seamier applications of online anonymity:

> People can anonymously transmit all sorts of illegal and injurious materials into public areas: copyright infringements, obscenity, stolen credit information, lies and slander, and so on. Individuals with a bone to pick against anyone else can get their licks in without fear of reprisal. Anonymous remailers are great for cowards. People who want to spread messages of hate and misunderstanding, but are unwilling to stand behind their views in public, can operate behind a wall of complete anonymity and inject a strong dose of thought pollution into the public arena.[132]

Another argument supporting disclosure of the origins of speech is quoted by Froomkin, ironically from an anonymous author: "'Disclosure advances the search for truth,' because when propaganda is anonymous it 'makes it more difficult to identify the self interest or bias underlying an argument.' "[133] Libel on the Internet is particularly pernicious, since once anything has been circulated via the 'Net, it becomes impossible to unring the bell of falsity.

The resurgence of hoaxes and rumors such as the Good Times virus[134] and the pathetic stories of Craig Shergold[135] and Jessica Mydek[136] illustrate the persistence of

undated, unsigned, and untrue messages in cyberspace. These unfounded, exaggerated, or obsolete stories, threats, and appeals circulate endlessly among the gullible on the 'Net. There is no reason to suppose they will ever stop.

One of the significant lessons from email hoaxes and chain letters is that unsigned, undated correspondence is always to be treated with skepticism. This principle of devaluing anonymous or pseudonymous communications will be used later in this chapter in a model for categorizing and sequestering communications as a function of their traceability.

Another growing pseudonymity issue is the use of substitute students in online degree-granting programs. One type of scam involves hiring students to pretend to be online students so that the criminals running the scheme can collect student-aid money from governments. In one massive scam, an Arizona women used imposters to steal $539,000 in student aid from the U.S. Government.[137] The availability of online courses and degrees has raised serious questions about the use of imposters who take exams on behalf of incompetent but rich candidates.[138]

70.5 SYSTEMS ANALYSIS OF ANONYMITY. Why can the usual protections of the criminal and civil law not deal with anonymous and pseudonymous communications? This problem is addressed by David Post and David Johnson.

Post and Johnson argue that geographically defined nation-states cannot reasonably cope with a virtual, boundaryless communications medium.[139] In the real world, geographical clustering combines with basic concepts of consent of the governed through some form of representation to legitimize the exercise of state power. Without the consent of the governed, state power fades insensibly into state tyranny. In their requirements analysis of possible systems of governance of cyberspace, they add that wherever possible, those affected by the conduct to be regulated have some say in framing the regulations.

However, in cyberspace, argue Post and Johnson, there is no geographical clustering. There is no "here" or "there" in cyberspace. "Location is indeterminate because there is no necessary relationship between electronic addressing … and the location of the addressee (machine or user) in physical space."

Post and Johnson applied the work of the scientist Stuart Kauffman on self-organizing systems to study the nature of rule-making in a complex system that can model the interactions among users in cyberspace.[140] Research on self-organization of complex systems suggests that optimum configurations of constraints (regulations) can evolve when there is some degree of aggregation (they refer to "patches") in the population. These aggregates represent groups where individuals sacrifice some of their preferences in return for overall improvement in the way the whole society works. Members of a patch share responsibility for governing their behavior.

One of the most striking findings of Post's and Johnson's research is that systems where most of the effects of an individual's actions are felt by others, outside its decision-making unit, lead to chaos or to suboptimal configurations of rules. Contrariwise, a balance between bringing consequences to bear on individuals in a "patch" and allowing effects to propagate through the larger population leads to more optimal results in the system.

As a result of the experiments, Post and Johnson suggest that one of the most powerful tools for rebuilding comity in cyberspace is grouping users by their Internet Service Providers. Each ISP can develop its own rules governing the behavior of members; sanctions for transgression of these local rules would include banishment from the ISP.

The authors examine the case of spam. As long as each ISP enforces technical measures against allowing fraudulent origination addresses, everyone in cyberspace can decide whether to filter out messages from any given ISP or not. ISPs that allow behavior judged harmful by others will limit the range of communication of their members.

Those that are viewed as restrictive will self-select their own members accordingly. Thus, without any global legislation, simply allowing individuals to choose ISPs that have published rules they like could lead to an effective self-regulation of communications. In essence, loudmouthed rumormongers would end up talking only to each other; junk email could be identified simply from its provenance; and even copyright violations could be punished by collective banning of communications from the offending ISPs.

Such a model is an instance of the ideal market of ideas in that objectionable ideas are not forbidden, they are just ignored. Of course, admirable and desirable ideas may also be ignored, but at least there is a choice involved. Access for communication becomes a form of currency in such a model—perhaps appropriate for the governance of cyberspace, the realm of electronic communications.

70.6 IMPLICATIONS AND DISCUSSION. Any attempt to restrict anonymity on the Internet would inevitably affect pseudonymity as well.[141] Anonymity and pseudonymity make law enforcement difficult enough, but these difficulties are exacerbated by the jurisdictional problems caused by a thoroughly decentralized communications medium like the Internet; "[W]ho should be setting the rules that apply to this new global medium?"[142] What, then, are some of the practical measures we can individually and collectively take to preserve the benefits of anonymity and pseudonymity without suffering the consequences of abuse?

70.6.1 Individuals, Families, and Schools. Absorbing and applying normative behavior begins in earliest childhood and continues throughout the development of the child's capacity for rationality and ethical judgment. Children should be taught that anonymity and pseudonymity are not acceptable under normal circumstances. The same methods that parents, teachers, and other adults use to teach children a visceral dislike of anti-social behaviors such as lying, cheating, stealing, and bullying should be applied to behavior in cyberspace.

It takes time to integrate morality into our technological universe. Twenty years ago, many drivers felt that driving under the influence of alcohol was adventurous. Today most people feel that it is stupid and irresponsible. Smoking in public is becoming rare. Many of us in northern cities have witnessed exiled smokers huddled together in the cold outside buildings where they once lit up with impunity.

Similarly, we need a consensus on good behavior in cyberspace.

Criminal hackers who break into computer systems and roam through users' private files should be viewed as peeping Toms. Criminals using computers to steal services should be recognized as thieves. Those who destroy records, leave logic bombs, and write viruses should be viewed as vandals. Hackers who smear obscenities in source code should be seen as twisted personalities in need of punishment and therapy. Government agencies proposing to interfere in electronic communications should be subject to scrutiny and intense lobbying.

Government-sponsored cyberespionage/warfare has become a significant threat to all. With current trends in government-sponsored cyberspying, many government agencies and corporations are at risk of losing very hard won and expensive technology

secrets to other countries who would rather anonymously and surreptitiously steal industrial secrets rather than to create the kind of economy that bears the fruits of innovation and technological earned advancement.

In a CNBC article entitled "Chinese Army Unit Is Seen as Tied to Hacking Against US" published on February 18, 2013, authors David E. Sanger, David Barboza, and Nicole Perlroth deliver a fascinating look at the Chinese Government's reported sponsorship of cyberspying.

> Unit 61398—formally, the 2nd Bureau of the People's Liberation Army's General Staff Department's 3rd Department—exists almost nowhere in official Chinese military descriptions. Yet intelligence analysts who have studied the group say it is the central element of Chinese computer espionage. The unit was described in 2011 as the "premier entity targeting the United States and Canada, most likely focusing on political, economic, and military-related intelligence" by the Project 2049 Institute, a nongovernmental organization in Virginia that studies security and policy issues in Asia.

The article contains specific information discovered by the security firm Mandiant; however, in the ultimate of ironies, several of the Chinese hackers (identified by nicknames) themselves use their very astute hacking skills to circumvent Chinese Government restrictions on access to social media sites:

> Mandiant discovered several cases in which attackers logged into their Facebook and Twitter accounts to get around China's firewall that blocks ordinary citizen's access … an embarrassing inconsistency to be sure.

Beyond such prohibitions and inhibitions of taboos, cyberspace needs the electronic equivalent of Emily Post. We need to discuss the immorality of virus writing, the ethical implications of logic bombs, and the criminality of electronic trespassing. We should teach children how to be good citizens of cyberspace, and not just in schools. We should sit down with computer-using youngsters and follow them through their adventures in cyberspace. Parents should ask their teenage whiz kids about hacking, viruses, software theft, and telephone fraud. We must bring the perspective and guidance of adult generations to bear on a world that is evolving faster than most of us can imagine.

The adolescent confraternity of criminal hackers and virus writers has already begun developing totems: the personae of Dark Avenger and Acid Phreak loomed over youngsters in the 1990s much as Robin Hood once did for other generations.

What we need now are taboos to match the totems.[143]

70.6.2 Ethical Principles. How do people make judgments about a course of action when there are no explicit guidelines? There are several kinds of principles that people use in reasoning about a new situation (see Chapter 43 of this *Handbook* for a more extensive discussion of ethical decision making).

70.6.2.1 Rights and Duties (Deontology). The concepts of rights ("Something that is due to a person or governmental body by law, tradition, or nature") and duties ("An act or a course of action that is required of one by position, social custom, law, or religion," according to the 1992 edition of the *American Heritage Dictionary*) should influence one's decisions.

Are any rights abridged by anonymity? For example, the right to know the source of a warning so that we may judge the motives and credibility of the statement is infringed when such a message is posted anonymously or pseudonymously.

Personal duties at issue in a decision on posting anonymous warnings about a competitor's product include notions of trust, integrity, and truthfulness, all of which are violated by such an act. We lower the trust of technical advice when we use anonymous postings; we damage the integrity of an entire profession; and we implicitly betray truthfulness by failing to identify the source of such information.

Professional duties or responsibilities also apply. We are expected to maintain appropriate professional relationships, but anonymous posting does not further a Web of trust among colleagues. Posting anonymous messages casts doubt on the goodwill of all the innocent people who are perceived as possibly being the author of such messages. In terms of maintaining efficacy, anonymous postings reduce the flow of information among professionals by aborting the possibility of private communication with the authors of the anonymous messages.

70.6.2.2 Consequentialism (Teleology). One approach to evaluating the ethical dimensions of a proposed act is to look at the possible consequences of the act. Does the action minimize actual and potential harm? Egoism looks at what is good for me or does the least harm to me. Anonymous posting of critical information about a competitor's product offers the potential of benefits to one's employer with minimal direct consequences. However, such behavior opens up the organization to less obvious consequences, such as the risk of blackmail, degradation of trust within the group, lowered morale, and departure of employees whose moral sensibilities are outraged by what they see as unethical behavior.

Utilitarianism views decisions in terms of what is good for the group or does the least harm for the group. The question here is the inclusivity of the "group." If the group includes all users and potential users of the defective product, then posting the information serves a good purpose; the decision on whether to post anonymously resolves to the same questions as those raised in discussions of the consequences for an organization of having an employee post messages anonymously. The climate of trust, the credibility of warnings in general, and respect for the entire industry can be harmed by anonymous postings of warnings.

An altruistic approach to decisions accepts that what is good for all may be worth doing even if there is some harm to one's self. By this yardstick, posting a warning with full attribution is definitely the preferred way of communicating information about a problem. Another altruistic approach, however, would be to inform the competitor of the flaw in its product via private communications. The hope here is that such altruism will be reciprocated and that the industry as a whole can benefit from improvement in all products.

70.6.2.3 Kant's Categorical Imperative. At a different level, Immanuel Kant's principles for judging the ethical dimensions of an act are immensely useful in all aspects of life. The principle of consistency asks, "What if everyone acted this way?" In our example, if everyone posted anonymous warnings, the credibility of all warnings would be damaged. In the absence of a mechanism of redress, unscrupulous people would contaminate the alerts with false information, making it difficult to trust any warning. Since some people would retaliate for fraudulent postings about their products by posting equally fraudulent attacks on their competitors, the system of alerts would collapse, causing harm to users and to producers.

Another principle Kant enunciated was that of respect: Are people treated as ends rather than means? The author of an anonymous message may not be thinking about the people affected by that message; they remain ciphers—amorphous, unknown entities

of no importance. In contrast, an empathic and ethical person remembers that every group consists of individual human beings with pretty much the same range of feelings as anyone else. Using them as a means of increasing market share is not respectful.

70.6.3 Corporations and Other Organizations. How can organizations contribute to the reduction of harmful anonymous or pseudonymous communications?

Every corporation, government department, nonprofit organization, educational institution, healthcare facility, banking or financial services organization, and so on should explicitly address the question of anonymity and forgery on the 'Net. Some practical guidelines to consider in setting organizational policy follow.

- No user of a corporate user ID should ever forge headers or use pseudonyms when communicating using corporate resources.
- Corporate firewalls should be configured to prevent all TCP/IP packets with forged origination IP addresses from passing outward through the firewall.
- SMTP servers should be configured to prevent any email from leaving a site with forged headers.
- All corporate email outbound from a site should be signed digitally by its author(s) to provide a basis, when necessary, for repudiation of unsigned and fraudulent email.

In addition, discussions of the ethical framework for making decisions about the use of technology should be integrated into employee training at all levels. Managers, in particular, ought to be cognizant of the principles of ethical decision making so that they can fluently guide their staff in practical problem solving.

As blogging, Twitter, and social networks (Facebook, LinkedIn, and others) allow for rapid and unmanaged corporate communications, written corporate policy on the use of blogging, Twitter, and social networking is a must.

Although sites like LinkedIn would traditionally not be anonymous, pseudo-personages are easy to construct and can be used for nefarious purposes. It does not take much of a stretch of the imagination to imagine a nefarious person/government constructing a series of false personas on LinkedIn and using those personas as a base from which to launch surveillance and information attacks on corporations through their employees who proudly proclaim their detailed information on LinkedIn.

70.6.4 Internet Service Providers. Because many users send email or post to Usenet groups, which with the growth of social media have seen a decline in use as social media sites have grown. Regardless, traffic still must traverse the Internet through accounts with ISPs, and these services have an obligation to become involved in preventing abuses of anonymity and pseudonymity. Except for some free services that are independently funded or that derive revenue from advertising rather than from user fees, ISPs must establish a relationship with their users in order to be paid; most use credit card accounts for this purpose. This method of billing inherently provides a link to the real-world identity of a user through the credit card issuers. Thus, for most ISPs, it is possible to enforce traceable anonymity and pseudonymity. Faced with a warrant or a subpoena duly issued by a judge, for example, most ISPs will be able to provide authorities with all the information needed to track down the person whose account was used to send a message. Linking a specific human being to the computer session used for sending that message is another matter.

70.6.5 A Free Market Model for Identity in Cyberspace. Some people use Caller-ID to refuse to accept phone calls from callers who mask their identities. Anyone who chooses not to be identified at the time of calling can simply leave a message.

What if users of ISPs could selectively reject communications from anonymous or unauthenticated pseudonymous sources? Email clients already allow users to reject or accept email from different top-level domains and to selectively block messages from specific sources known to generate enormous amounts of spam. The Chinese numerical domains of the form *nnn*.com (e.g., 123.com, 163.com), for example, are notorious for spewing spam messages written entirely in Chinese or which, using peculiar approximations to English, advertise concrete sewage pipes to information assurance professors; these nonsensical advertisements are delivered all over the world, probably due to the activities of swindlers who trick gullible neophytes into paying for supposedly targeted email campaigns.[144] Blocking all traffic from these domains is easy.

What might happen over time as a result of allowing communications only with selected ISPs, depending on their terms of service? Eventually, there would likely be an equilibrium in which those users who wished to send and receive anonymous, untraceable email could subscribe to those ISPs supporting that kind of communication. Others could automatically block email from unwanted ISPs. Furthermore, individuals who wished to use anonymity or pseudonymity sometimes could subscribe to more than one ISP and take advantage of their different policies. Computer systems managers who wanted to deal only with other professionals at work could use a restrictive ISP; however, they also could use a different ISP to post and read messages in a highly charged, free-wheeling discussion group about the politics of gun control, without having to reveal their real names or fear that they could ultimately be traced through a pseudonym.

As the forces of the marketplace continued to work on ISPs, there might be further evolution toward different degrees and types of communication blockage. ISPs with a reputation for harboring miscreants who libel and defame others without cause could find themselves being shut out of an increasing number of reputable communities of users. Those whose users exercised moderation and responsibility might find themselves being received by a widening circle of ISPs.

The advantage of this model is that individuals could exert power over their electronic environment by voting with their subscriptions, but no one would be censored by bureaucrats or tyrants. Just as Hyde Park in England allows both geniuses and crackpots to speak yet forces no one to listen, the electronic Hyde Park would provide a mechanism for shutting out the lunatics while letting the loonies talk to each other as they wish.

What this model would not permit is the imposition of unwanted communications on powerless victims. This model would spell the end of unsolicited commercial email. If digital signatures indicating the source ISP from which executable code was first distributed became commonplace, the same mechanism could seriously interfere with the distribution of viruses and other harmful programs. ISPs that became known for harboring virus writers or distributors would see their credibility as communications partners eroded. This model does not require individuals to sign their product or messages; the only constraint is that the ISP do so.

What about individuals who own their own domain on the ‘Net? The same principles would apply. The modified SMTP software would automatically sign all output from their sites. Any site that refused to provide digital signatures could be excluded by any ISP that chose to apply such an exclusionary rule. Even postings to blogs

and other discussion groups on the Web could be constrained by verification of the reliability of identities associated with the originators; those insisting on being anonymous or pseudonymous would have the freedom not to post on the sites demanding authentication.

70.6.6 Governments. The role of governments in cyberspace is complex. On one hand, several governments—notably that of the United States—have contributed to the development of the Internet through legislation and funding. On the other hand, many governments, especially totalitarian regimes, are intolerant of unfettered communications. Under the communist regimes, most countries in the Soviet Union and its satellites made ownership of unregistered spirit duplicators, photocopiers, fax machines, and modems illegal. Today, China, in the grip of a tyrannical regime, behaves equally rigidly with respect to electronic communications.

No matter how carefully crafted they are, attempts to apply legal constraints on anonymity in cyberspace will be undermined by inconsistent government regimes throughout the globe. The least restrictive geographical entities will subvert more restrictive jurisdictions.[145] If, say, Restrictopolis were to impose strictures on anonymous Internet use, its anonymity-seeking citizens might be able to use the ISPs in Liberalopolis, where the rules would be much freer.

On a less practical, more philosophic level, there are profound objections to any government regulation of the Internet. As Lewis writes: "We do not outlaw wig shops or Halloween masks just because some people use them for illegal or immoral purposes. We do not require caller-ID services for everyone just because some people make obscene or harassing phone calls. Nor should we strip the cloak of online anonymity from everyone, including those who legitimately need privacy, just to prevent sickos from abusing it."[146]

In the United States, the government would likely run up against strong constitutional guarantees of speech, especially political speech—and including anonymous political speech—if it tried to ban anonymity outright.[147] A particularly significant setback for government attempts to control anonymity in the United States came in June 1997. The case began in January 1997.[148] As Declan McCullagh described the judgment:

> In Georgia, Judge Marvin Shoob ruled that a state law forbidding anonymity online is unconstitutional since it violates free speech and free association rights. The law is so broadly written, the judge indicated, that even America Online screen names could be considered illegal. Judge Shoob "understood clearly the very strong need for our plaintiffs to communicate anonymously," the ACLU's Ann Beeson says. Both judges [in Georgia and in a similar case in New York] issued preliminary injunctions barring the state attorneys general from enforcing the laws. …
>
> Georgia's Judge Shoob, in a 21-page opinion, ruled that the law—that the Democrat-controlled legislature passed in haste last year to muzzle a dissident Republican representative—violated the First Amendment.
>
> This echoes a recent Supreme Court case, *McIntyre v. Ohio*, in which the justices ruled that the right to anonymity extends beyond political speech; that requiring the author's name on a leaflet is unconstitutional; that writing can be more effective if the speaker's identity is unknown.[149]

Democratic governments worldwide would do better to stay out of cyberspace and allow users to develop their own transnational solutions for governing behavior, including the use of anonymity and pseudonymity.

For a discussion of the explosive revelations of massive data collection by the National Security Agency of the United States, see Chapter 69 in this *Handbook.*

70.7 CONCLUDING REMARKS. Readers of this brief review of anonymity and pseudonymity are urged to consider these key observations and recommendations:

- Anonymous and pseudonymous communications are inherently associated with an increased incidence of antisocial behavior through deindividuation, social containment, general crime theory, and social learning.
- There are circumstances where anonymity and pseudonymity are useful tools in the defense of liberty and justice.
- Anonymous and pseudonymous electronic communications have already been used to harass victims, damage commercial interests, and launch hoaxes and rumors into cyberspace.
- The major problems of anonymity and pseudonymity in cyberspace can be avoided by the use of traceable identification.
- Making ISPs responsible for determining reasonable methods for enforcing their chosen level of strong identification and authentication will allow a nongovernmental, nonlegalistic approach to reducing abuse by anonymous and pseudonymous Internet users.
- All electronic communications ought to be tagged with unforgeable authenticators of identity.
- Individuals, families, and schools have a role to play in integrating cyberspace into the moral universe of children.
- Corporations and other organizations ought to integrate ethical decision making into their management procedures and corporate policies, and enforce those policies especially related to social media.
- Governments will continue to fail in their efforts to govern cyberspace because electronic communications networks are inherently divorced from geographical jurisdictions.
- A well-conceived, flexible strategy and resultant policies are essential as the instantaneous nature, growth, and evolution of social media and mobile computing exposes far more attack surface, which can be manipulated by anonymity and pseudoanonymity. Possessing the tools and vision to match the evolutionary pace of the Internet is key to managing the future.

70.8 SUMMARY. The growth of the Internet has increased the use of anonymity and pseudonymity in electronic communications. Internet users must be able to preserve the benefits of privacy while fighting the abuses of anonymous and pseudonymous people. In the real world, identity resides in the ways that individuals are recognized and held responsible for their actions; in cyberspace, identity is potentially just a user ID. Social psychologists have found that anonymity can contribute to deindividuation—a state of loss of self-awareness, lowered social inhibitions, and increased impulsivity.

This chapter suggests practical applications of these insights from social psychology for managers concerned with reducing abusive behavior by their own employees. In addition, the chapter addresses the wider problem that, given the social psychology of anonymity, abuses of the Internet are certain to continue. There must develop a collective response to incivility and irresponsibility, without falling into authoritarian strictures on speech. This chapter further suggests that a free market approach using accessibility to communications as a kind of currency may help the 'Net evolve toward

a more civil society. By blocking email from ISPs that fail to enforce acceptable standards for a given community of users, 'Net users can sort themselves into groups that tolerate or welcome different levels of anonymity and pseudonymity. No government intervention would be required under such a system.

In addition, this chapter suggests a framework for reaching into the early years of the educational system to help integrate cyberspace into the moral universe of children worldwide. In addition to being a moral imperative to support educational efforts on computer ethics, such programs are in the best economic interests of industry, academia, and government systems managers.

70.9 FURTHER READING

Bailey, M. *Complete Guide to Internet Privacy, Anonymity & Security*. Nerel Online, 2011.

Baggio, B., and Y. Beldarrain, eds. *Anonymity and Learning in Digitally Mediated Communications: Authenticity and Trust in Cyber Education*. IGI Global, 2011.

Boellstorff, T. *Coming of Age in Second Life: An Anthropologist Explores the Virtually Human*. Princeton University Press, 2008.

Eydie, A. *How to be Anonymous Online—A Quick Step-By-Step Manual*. Amazon Digital Services, 2013

Federrath, H., ed. *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, Berkeley, CA, July 25–26, 2000. Proceedings. Springer

Griffin, R. J. *The Faces of Anonymity: Anonymous and Pseudonymous Publication from the Sixteenth to the Twentieth Century*. New York: Palgrave Macmillan, 2003

Kjaerland, M. *Electronic Civil Disobedience. A Differentiation between "Hactivists" Based on Target and Message of Web Site Hacks*. Dissertation: The University of Liverpool, Department of Psychology, 2000.

Nicoll, C., J. E. J. Prins, and M. J. M. van Dellen, eds. *Digital Anonymity and the Law: Tensions and Dimensions*. Asser Press, 2003

Robbins, S., and M. Bell. *Second Life for Dummies*. Hoboken, NJ: John Wiley & Sons, 2008

Stryker, C. *Epic Win for Anonymous: How 4chan's Army Conquered the Web*. Overlook Hardcover, 2011

Stryker, C. *Hacking the Future: Privacy, Identity, and Anonymity on the Web*. Overlook Hardcover, 2012

Ziccardi, Giovanni. *Resistance, Liberation Technology and Human Rights in the Digital Age. Law, Governance and Technology Series*, Volume **7**, Pompeu Casanovas and Giovanni Sartor, eds. Springer Netherlands, 2013. http://cryptome.org/2013/03/hacking-digital-dissidence.pdf

70.10 NOTES

An asterisk (*) in notes marks a secondary reference to classic social psychology works drawn from D. G. Myers, *Social Psychology,* 4th ed. (New York: McGraw-Hill, 1993); R. A. Lippa, *Introduction to Social Psychology,* 2nd ed. (Pacific Grove, CA: Brooks/Cole Publishing, 1994); or D. O. Sears, L. A. Peplau, and S. E. Taylor, *Social Psychology,* 7th ed. (Englewood Cliffs, NJ: Prentice-Hall, 1991).

1. U.S. Department of Justice, United States Attorney's Office for the District of Maryland, "Identity Theft Working Group Increases Federal Prosecutions:

Maryland Crime Victims' Resource Center to Assist Victims," February 21, 2008, www.justice.gov/usao/md/news/archive/IdentityTheftWorkingGroupIncreasesFederalProsecutions.html

2. www.nytimes.com/2012/01/19/world/asia/china-expands-program-requiring-real-name-registration-online.html?_r=0&pagewanted=print
3. R. Abelson, "By the Water Cooler in Cyberspace, the Talk Turns Ugly," *New York Times*, www.nytimes.com/2001/04/29/technology/29HARA.html
4. R. D. Clifford, *Cybercrime: The Investigation, Prosecution and Defense of a Computer-Related Crime*, 3rd ed. (Carolina Academic Press, 2011)
5. J. Kornblum, "Net Crime Begs Questions: Who to Call?" C|Net News, August 5, 1997; http://news.cnet.com/2102-1023_3-202141.html
6. E. Kabay, "The INFOSEC Year in Review 1997", see www.mekabay.com/iyir/1997.pdf
7. WHO@ home page: www.haltabuse.org
8. CyberAngels home page: www.guardianangels.org/cyberangels.php
9. Matt Pearce, "Web Plays Incendiary Role in Ohio High School Rape Case," *Los Angeles Times*, January 7, 2013, http://articles.latimes.com/print/2013/jan/07/nation/la-na-nn-ohio-stuebenville-sexual-assault-20130106
10. J. Kornblum, "Antispam Efforts Heat Up," C|Net News, November 14, 1997, http://news.cnet.com/Antispam-efforts-heat-up/2100-1023_3-205363.html
11. Tracy Laquey Parker, Zilker Internet Park, Inc., Patrick Parker, Peter Rauch, Texas Internet Service Providers Association and EFF-Austin, Plaintiffs, v. C.N. Enterprises and Craig Nowak, Defendants, Final Judgement, No. 97-06273 (345th Judicial District of Texas, November 10, 1997), www.loundy.com/CASES/Parker_v_CN_Enterprises.html
12. Internet Corporation for Assigned Names and Numbers (ICANN). "Top-Level Domains (gTLDs)," 2013, www.icann.org/tlds
13. R. Clarke, "Human Identification in Information Systems: Management Challenges and Public Policy Issues," *Information Technology & People* 7, No. 4 (1994): 6–37, www.rogerclarke.com/DV/HumanID.html
14. A. M. Froomkin, "Anonymity and Its Enmities," *Journal of Online Law*, 1995, article 4, http://groups.csail.mit.edu/mac/classes/6.805/articles/anonymity/froomkin.html
15. OSHA, "Milliken & Company VPP Star Site." OSHA Voluntary Protection Programs, Success with VPP, 2006, www.osha.gov/dcsp/success_stories/vpp/reg4_ss_milliken_elm.html
16. Clarke, "Human Identification in Information Systems."
17. D. G. Post, "Knock Knock, Who's There?: Anonymity and Pseudonymity in Cyberspace."
18. Froomkin, "Anonymity and Its Enmities."
19. G. Le Bon, *The Crowd: A Study of the Popular Mind* (Macmillan, 1896).*
20. L. Festinger, A. Pepitone, and T. Newcomb, "Some Consequences of Deindividuation in a Group," *Journal of Abnormal and Social Psychology*, 47 (1952): 382–389.*
21. P. G. Zimbardo, "The Human Choice: Individuation, Reason and Order versus Deindividuation, Impulse, and Chaos," in *Nebraska Symposium on Motivation*, ed. W. J. Arnold and D. Levine (Lincoln: University of Nebraska Press, 1970).*

22. Montaldo, C. "The Growing Problem of Road Rage," About.com (2013), http://crime.about.com/od/victims/p/roadrage.htm
23. W. C. Reckless, *The Crime Problem*, Meredith Publishing Company, New York, 1967.
24. W. C. Reckless, *The Crime Problem*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1973.
25. C. H. McCaghy, T. A. Capron, J. D. Jamieson, and S. H. Carey, *Deviant Behavior: Crime, Conflict, and Interest Groups* (Boston: Pearson Education, Inc., 2006).
26. J. McGuire, *Understanding Psychology and Crime* (New York: Open University Press, 2004).
27. McCaghy et al., *Deviant Behavior.*
28. McGuire, *Understanding Psychology and Crime*.
29. R. W. Taylor, T. J. Caeti, D. K. Loper, E. J. Fritsch, and J. Liederbach, *Digital Crime and Digital Terrorism* (Upper Saddle River, New Jersey: Pearson Education, 2006).
30. R. L. Akers, *Criminological Theories: Introduction and Evaluation* (Los Angeles: Roxberry Publishing Company, 1994).
31. M. M. Lanier and S. Henry, *Essential Criminology* (Boulder, CO: Westview Press, 2004).
32. Akers, *Criminological Theories*.
33. McCaghy et al., *Deviant Behavior.*
34. McGuire, *Understanding Psychology and Crime*.
35. Taylor et al., *Digital Crime and Digital Terrorism*.
36. McGuire, *Understanding Psychology and Crime*.
37. Taylor et al., *Digital Crime and Digital Terrorism*.
38. Akers, *Criminological Theories*.
39. McCaghy et al., *Deviant Behavior.*
40. Taylor et al., *Digital Crime and Digital Terrorism*.
41. Taylor et al., *Digital Crime and Digital Terrorism*.
42. McGuire, *Understanding Psychology and Crime*.
43. Taylor et al., *Digital Crime and Digital Terrorism*.
44. Taylor et al., *Digital Crime and Digital Terrorism*.
45. McGuire, *Understanding Psychology and Crime*.
46. Akers, *Criminological Theories*.
47. R. L. Burgess and R. L. Akers, "A differential association-reinforcement theory of criminal behavior," *Social Problems* 14 (1966): 128–147.
48. Lanier and Henry, *Essential Criminology*.
49. McCaghy et al., *Deviant Behavior.*
50. McGuire, *Understanding Psychology and Crime*.
51. Akers, *Criminological Theories*.
52. S. Hinduja, *Music Piracy and Crime Theory* (New York: LFB Scholarly Publishing LLC, 2006).
53. R. L. Akers, *Social Learning and Social Structure: A General Theory of Crime and Deviance* (Boston: Northeastern University Press, 1998).

54. R. L. Matsueda, "The Current State of Differential Association Theory," *Crime and Delinquency* 34, no. 3 (1988): 277-306.
55. Akers, *Criminological Theories*.
56. Akers, *Social Learning and Social Structure*.
57. Akers, *Criminological Theories*.
58. Taylor et al., *Digital Crime and Digital Terrorism*.
59. McGuire, *Understanding Psychology and Crime*.
60. McCaghy et al., *Deviant Behavior.*
61. Akers, *Criminological Theories*.
62. R. White and F. Haines, *Crime and Criminology* (New York: Oxford University Press, 1996).
63. Lanier and Henry, *Essential Criminology*.
64. Matsueda, *The Current State of Differential Association Theory*.
65. K. Schuessler, *Edwin H. Sutherland on Analyzing Crime* (Chicago: The University of Chicago Press, 1973).
66. Matsueda, *The Current State of Differential Association Theory*.
67. C. S. Sellers, J. K. Cochran, and K. A. Branch, "Social Learning Theory and Partner Violence: A Research Note," *Deviant Behavior* 26, no. 4 (2005): 379–395.
68. Schuessler, *Edwin H. Sutherland on Analyzing Crime*.
69. Matsueda, *The Current State of Differential Association Theory*.
70. Sellers et al., "Social Learning Theory and Partner Violence."
71. Akers, *Social Learning and Social Structure*.
72. R. L. Akers and G. F. Jensen, eds., *Advances in Criminological Theory,* vol. 11, *Social Learning Theory and the Explanation of Crime: A Guide for the New Century* (New Brunswick, NJ: Transaction Publishers, 2003).
73. T. Thornberry, ed., *Advances in Criminological Theory,* vol. 7, *Developmental Theories of Crime and Delinquency* (New Brunswick, NJ: Transaction Publishers, 1997).
74. Matsueda, *The Current State of Differential Association Theory*.
75. Matsueda, *The Current State of Differential Association Theory*.
76. Zimbardo, "The Human Choice."
77. R. I. Watson, Jr., "Investigation into Deindividuation Using a Cross-Cultural Survey Technique," *Journal of Personality and Social Psychology* 25 (1973): 342–345.*
78. J. H. Kerr, *Understanding Soccer Hooliganism* (Buckingham: Open University Press, 1994).*
79. B. Buford, *Among the Thugs* (Mandarin Paperbacks, 1991). Also available via Kindle.
80. E. Diener, S. C. Fraser, A. L. Beaman, and R. T. Kelem, "Effects of Deindividuating Variables on Stealing by Halloween Trick-or-Treaters," *Journal of Personality and Social Psychology* 33 (1976): 178–183.*
81. S. Prentice-Dunn and R. W. Rogers, "Effects of Deindividuating Situation Cues and Aggressive Models on Subjective Deindividuation and Aggression," *Journal of Personality and Social Psychology* 39 (1980): 104–113*; and E. Diener,

"Deindividualtion, Self-Awareness, and Disinhibition," *Journal of Personality and Social Psychology* 37 (1979): 1160–1171.*

82. M. Csikszentmihalyi, *Flow: The Psychology of Optimal Experience* (Harper Perennial Modern Classics, 1990, republished 2008).
83. A. Cooper, M. P. Safir, and A. Rosenmann, "Workplace Worries: A Preliminary Look at Online Sexual Activities at the Office—Emerging Issues for Clinicians and Employers," *CyberPsychology & Behavior* 9 (2006): 22–29.
84. C. B. Spivey and S. Prentice-Dunn, "Assessing the Directionality of Deindividuation: Effects of Deindividuation, Modeling, and Private Self-Consciousness on Aggressive and Prosocial Responses," *Basic and Applied Social Psychology* 11 (1990): 387–403.*
85. K. J. Gergen, M. M. Gergen, and W. H. Barton, "Deviance in the Dark," *Psychology Today* (October 1973): 129–130.*
86. L. Detweiler, "Anonymity on the Internet FAQ" (1993), www.websteruniv.edu/~bumbaugh/net/anonfaq.html*;* and L. Detweiler, "The Joy of Pseudospoofing" (1993), *no longer available online;* but see L. Detweiler, "The Snakes of Medusa and Cyberspace: Internet Identity Subversion" *The Risks Digest* 15, no. 25 (November 10, 1993), http://catless.ncl.ac.uk/Risks/15.25.html#subj4
87. Detweiler, "Anonymity on the Internet FAQ."
88. R. Clarke, "Introduction to Dataveillance and Information Privacy, and Definitions of Terms," Roger Clarke's Web-Site, 1999, www.rogerclarke.com/DV/Intro.html
89. Clarke, "Human Identification in Information Systems."
90. Detweiler, "Joy of Pseudospoofing."
91. Detweiler, "Joy of Pseudospoofing."
92. Renner, T., and M. E. Kabay. "Dishonest Methods Involving Online Reputation Management," *Network World | Security Strategies Alert*, August 10, 2011, www.networkworld.com/newsletters/sec/2011/080811sec2.html
93. For current resources on electronic voting issues, see the Electronic Privacy Information Center's files at http://epic.org/privacy/voting/
94. Froomkin, "Anonymity and Its Enmities."
95. Froomkin, "Anonymity and Its Enmities."
96. C. Johnson, "Anonymity On-line? It Depends on Who's Asking," *Wall Street Journal*, November 24, 1995, p. B1.
97. L. Lessig, D. Post, and E. Volokh, "Cyberspace Law for Non-Lawyers. Lesson 23—Privacy 11: Privacy: Self-Help: Anonymity, Part 1," Originally available via the Social Science Research Network (SSRN) Website, 1997. Now available via the Internet Archive Wayback Machine: http://web.archive.org/web/19970416041146/http://www.ssrn.com/cyberlaw/
98. L. Lessig, D. Post, and E. Volokh, "Cyberspace Law for Non-Lawyers. Lesson 24—Privacy 12: Privacy: Self-Help: Anonymity, Part 2" Originally available via the Social Science Research Network (SSRN) Website, 1997. Now available via the Internet Archive Wayback Machine: http://web.archive.org/web/19970416041146/http://www.ssrn.com/cyberlaw/
99. Clarke, "Introduction to Dataveillance and Information Privacy."
100. E. Mahyuni, "Gone in 6 Seconds: Hotmail CAPTCHA Hacked by Spammers," *Tech.Blorge*, April 16, 2008, http://tech.blorge.com/Structure:%20/2008/04/16/gone-in-6-seconds-hotmail-captcha-hacked-by-spammers

101. C. Pappas, "The A to Z of Internet Sleaze," *Home Office Computing* 15, no. 8 (1997):70.
102. Detweiler, "Anonymity on the Internet FAQ."
103. W. Rodger, "Zero-Knowledge to Close Anonymity Service," SecurityFocus, 2001, www.securityfocus.com/news/262
104. K. Zetter, "Rogue Nodes Turn Tor Anonymizer into Eavesdropper's Paradise," Wired, September 10, 2007, www.wired.com/politics/security/news/2007/09/embassy_hacks?currentPage=all
105. "Maimonides' Eight Levels of Charity." Mishneh Torah, Laws of Charity, 10:7–14. Chabad.org Library, www.chabad.org/library/article_cdo/aid/45907/jewish/Eight-Levels-of-Charity.htm
106. Electronic Frontier Foundation, "Anonymous Critics' Identities Protected," Electronic Frontier Foundation Website, August 26, 2005, https://www.eff.org/press/archives/2005/08/26
107. Clarke, "Introduction to Dataveillance and Information Privacy."
108. Froomkin, "Anonymity and Its Enmities"; and A. M. Froomkin, "Flood Control on the Information Ocean: Living with Anonymity, Digital Cash and Distributed Databases," *University of Pittsburgh Journal of Law and Commerce* 15 (1996): 395, http://osaka.law.miami.edu/~froomkin/articles/ocean1.htm
109. J. Ubois, "Anonymity Has Its Place," *MIDRANGE Systems* 8, no. 8 (1995): 28.
110. Detweiler, "Anonymity on the Internet FAQ."
111. P. Steiner, "On the Internet, Nobody Knows You're a Dog." *New Yorker*, July 5, 1993, http://i2.kym-cdn.com/entries/icons/original/000/011/625/Internet_dog.jpg
112. B. Paisley, "Online" (2008), YouTube video of performance in Nashville: www.youtube.com/watch?v=UE6iAjEv9dQ
113. Detweiler, "Anonymity on the Internet FAQ."
114. Froomkin, "Flood Control on the Information Ocean."
115. Froomkin, "Flood Control on the Information Ocean."
116. D. Post, "Knock Knock, Who's There?: Anonymity and Pseudonymity in Cyberspace," 1995 [article not available online in June 2013, according to author].
117. G. Stringer, "Virtual Communities," course notes, "The Internet," Creative Media and Information Technology course MIT 2114/2214, 2005. University of Exeter (UK), Virtual Communities, http://services.exeter.ac.uk/cmit/modules/the_internet/webct/ch-communities.html
118. Ellen Gorsevski, "Running Free in Angelina Jolie's Virtual Body: The Myth of the New Frontier and Gender Liberation in Second Life," Chapter 8 in *Millennial Mythmaking: Essays on the Power of Science Fiction and Fantasy Literature, Films and Games*, John R. Perlich and David Whitt, eds. McFarland, 2009.
119. Clarke, "Introduction to Dataveillance and Information Privacy."
120. Detweiler, "Anonymity on the Internet FAQ."
121. Froomkin, "Anonymity and Its Enmities."
122. Froomkin, "Flood Control on the Information Ocean."
123. EPIC. "USA Patriot Act." Electronic Privacy Information Center, 2013, http://epic.org/privacy/terrorism/usapatriot

124. Brett Max Kaufman, "ACLU Files Lawsuit Challenging NSA's Patriot Act Phone Surveillance." ACLU Website, June 11, 2013, www.aclu.org/blog/national-security-technology-and-liberty/aclu-files-lawsuit-challenging-nsas-patriot-act-phone
125. The Tor Project, https://www.torproject.org
126. InfraGard home page: www.infragard.net
127. Detweiler, "Anonymity on the Internet FAQ."
128. Froomkin, "Anonymity and Its Enmities."
129. Froomkin, "Anonymity and Its Enmities."
130. Post, "Knock Knock, Who's There?"
131. Detweiler, "Anonymity on the Internet FAQ."
132. L. J. Rose, *NetLaw: Your Rights in the Online World* (New York: Osborne/McGraw-Hill, 1994), pp. 183–184.
133. Froomkin, "Flood Control in the Information Ocean."
134. Good Times Virus Hoax FAQ: www.cityscope.net/hoax1.html; but see also the satirical version at http://ifaq.wap.org/computers/goodtimesfaq.html
135. B. Mikkelson, "Craig Shergold," Snopes.com (2007): www.snopes.com/inboxer/medical/shergold.asp
136. B. Mikkelson, "American Cancer Society Hoax." Snopes.com (2006), www.snopes.com/inboxer/medical/cancer.asp
137. M. Parry, "Online Scheme Highlights Fears About Distance-Education Fraud," *Chronicle of Higher Education,* January 13, 2010, http://chronicle.com/article/Online-Scheme-Highlights-Fears/63517
138. M. E. Kabay, "Academic Dishonesty in Online Coursework," Presentation at 2012 eCampus Security Conference, Dartmouth College, http://www.ists.dartmouth.edu/docs/ecampus/2012/2012ecampus_kabay.pdf
139. D. Post and D. R. Johnson, "The New 'Civic Virtue' of the Internet: Lessons from Models of Complex Systems" (1997), http://journals.uic.edu/ojs/index.php/fm/rt/printerFriendly/570/491
140. Post and Johnson, "The New 'Civic Virtue' of the Internet"; and S. A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution* (Oxford University Press, 1993).
141. Post, "Knock Knock, Who's There?"
142. Post and Johnson, "The New 'Civic Virtue' of the Internet."
143. M. E. Kabay, "Totem and Taboo in Cyberspace: Integrating Cyberspace into Our Moral Universe," Proceedings of the 17th National Computer Security Symposium, Baltimore, MD, October 11–14, 1994. Reprinted in NCSA News (June 1995): 4, online at www.mekabay.com//ethics/totem_taboo_cyber.pdf
144. M. E. Kabay, "Incomprensible Spam: Growing Global Economices Falling Prey to Scammers." *NetworkWorld | Security Strategies*, June 6, 2011, www.networkworld.com/newsletters/sec/2011/060611sec1.html
145. Froomkin, "Anonymity and Its Enmities."
146. P. Lewis, "Cloaks and Daggers: Online Anonymity Is a Blessing and a Curse," *Home Office Computing* 14(7):133 (1996).
147. Lessig et al., "Cyberspace Law for Non-Lawyers, Lesson 24."

148. D. McCullagh, "Brick by Brick," *The Netly News* (editorial), 1997, http://tech.groups.yahoo.com/group/cypherpunks/message/973?o=1&d=-1

149. D. McCullagh, "Courts Strike Down New York and Georgia Net-Censorship Laws," sent Friday, June, 1997, 13:49:06 -0700 (PDT), on fight-censorship-announce@vorlon.mit.edu mailing list. www.well.com/~declan/nym/georgia.0797.txt

CHAPTER 71

HEALTHCARE SECURITY AND PRIVACY

Paul Brusil

71.1 OVERVIEW/SYNOPSIS 71·2

71.2 MOTIVATION FOR HEALTHCARE INFORMATION PRIVACY AND SECURITY 71·4

71.3 HEALTHCARE SECTOR THREATS AND VULNERABILITIES LANDSCAPE 71·4

71.3.1 Healthcare Data Breaches are Growing 71·10
71.3.2 Scope of Attack Vectors and Breach Consequences Is Large 71·10
71.3.3 More Attack Vectors are Emerging 71·12
71.3.4 The Bottom Line 71·13

71.4 HEALTHCARE INFORMATION PROTECTION REQUIREMENTS IN U.S. LAWS 71·14

71.5 THE HIPAA AND HITECH HEALTHCARE INFORMATION PROTECTION REGULATIONS 71·16

71.5.1 HIPAA's Privacy and Security Fundamental Principles 71·18
71.5.2 HIPAA Privacy Rule 71·19
71.5.3 HIPAA Security Rule 71·21
71.5.4 HIPAA Breach Notification Interim Rule 71·24
71.5.5 HIPAA Protection Rules Enforcement, Penalties, and Liabilities 71·27
71.5.6 HITECH's HIPAA Update 71·28
71.5.7 Demonstrating HIPAA Compliance 71·31
71.5.8 Related Federal Guidelines 71·34
71.5.9 Comparing Federal and State Healthcare-Relevant Security and Privacy Laws 71·34

71.6 OTHER PUBLIC SECTOR HEALTHCARE PRIVACY AND SECURITY INITIATIVES 71·40

71.6.1 The HITECH Act: Moving to Electronic Health Information 71·41
71.6.2 Healthcare Information Protection in the U.S. DoD Healthcare Environment 71·45
71.6.3 Evolving and Emerging Government Healthcare Information Protection Initiatives 71·47

71.7 PRIVATE SECTOR HEALTHCARE PRIVACY AND SECURITY INITIATIVES 71·48
71.7.1 HITRUST 71·48
71.7.2 Medical Devices Security Initiatives 71·50

71.8 SECURITY AND PRIVACY IMPACTS OF TECHNOLOGY TRENDS 71·51
71.8.1 Mobile and Personal Technology 71·52
71.8.2 Virtualized Technology 71·53
71.8.3 Personalized Medical Technology 71·53
71.8.4 Big Data Trends 71·53
71.8.5 Other Factors Associated with Technology Trends 71·54

71.9 SUMMARY 71·55

71.10 FURTHER READING 71·56

71.11 NOTES 71·56

71.1 OVERVIEW/SYNOPSIS. Healthcare has become big business. In the United States it has grown to represent a very large fraction of the domestic economic output. The delivery of healthcare depends on use of personal healthcare and personal financial information. Threats to such information grow as information is aggregated into and used by various healthcare entities. Motives for compromising private healthcare-related information are evolving from sneaking celebrities' medical information to more serious, consequential, profit-driven motives such as medical identity theft.

Trust in the healthcare system is essential. Initially, the demands on protecting healthcare-relevant information were driven by simple, government-created, generalized regulations on privacy. They have since grown to significant, far-reaching, federal, state, and industry security and privacy regulations specific to both the healthcare-provisioning side and the business processing side of the healthcare system. Furthermore, as the healthcare system evolves to the use of always newer information technology (IT), IT service delivery paradigms and healthcare service delivery strategies, maintaining trust may not always be simple or straightforward.

The healthcare sector move to adopt IT is providing tectonic changes to the healthcare industry. As IT is embraced, the drive to lower healthcare costs will accelerate. A new healthcare business will capitalize on innovative use of IT. Big changes are in the offing as healthcare adopts technology.

First up is the current push to move to electronic patient records—so-called electronic health records (EHRs)—that will power the new digital healthcare sector. Next up is the electronic sharing of such EHRs and the personal healthcare data they contain among healthcare providers, specialists, and patients, in part via health information exchanges. Electronic data sharing will include real-time patient monitoring and patient/provider messaging.

Following the goal of electronic healthcare data sharing and the near-term focus on ensuring that healthcare IT meets the frantic pace of newly instituted government mandates, attention is expected to turn to adopting yet other new technologies. These technologies will offer the potential for big breakthroughs in personalized patient care outcomes and wellness. Analytics technologies[1] such as predictive analytics and big data analytics, as well as decision-support systems, expect to be key.

Such technologies will be essential in developing advanced medical knowledge from aggregated and mined healthcare data. Such knowledge can improve overall healthcare system quality and cost. On a case-by-case basis, such knowledge can be integrated with real-time, patient-specific data to diagnose and to treat individual patients in more efficacious ways. Today's advanced technologies, and other technologies yet to

be conceived, will likely be slow in being incorporated across the board into routine medical workflows; provider trust in such technologies and associated ROI potentials will need to be sufficiently established.

As such, the ongoing and emerging dramatic shifts to reliance on new underlying IT is forcing equally significant attention to the security and privacy needs of the personal data fueling the healthcare sector's emerging digital industry. After all, lack of trust in the accuracy and completeness of shared personal healthcare information can jeopardize the delivery of quality care and have devastating consequences in the new healthcare system model, which depends on secure IT for the coordination of care and the advancement of healthcare knowledge.

The goal of this chapter is to provide a sense of the existing and emerging security and privacy issues facing the healthcare sector during the middle of the second decade of the 2000s.

It examines many of the more significant regulations, guidelines, and initiatives that progress the implementation and secure use of protected information to provide and to administer care in a healthcare environment.[2] References and links are provided for readers to learn more in-depth and evolving details themselves.

This chapter begins by summarizing the motivation for security and privacy in the handling of healthcare-related information. In Section 71.3 it provides a picture of the existing and emerging threats and vulnerabilities that exist in the healthcare sector. The requirements for healthcare sector privacy and security that have arisen from state laws and early Federal laws are identified in Section 71.4.

The significant and comprehensive HIPAA (Health Insurance Portability and Accountability Act of 1996) healthcare information protection "Rules" and follow-on HITECH (Health Information Technology for Economic and Clinical Health Act of 2009) modifications to these rules that emerged over the past 17 years are presented in Section 71.5. Initiatives related to demonstrating HIPAA compliance are considered in Section 71.5.7. A tabulation of government guidelines and emerging government initiatives that pertain to HIPAA and other aspects of protecting healthcare information is given in Section 71.5.8. Section 71.5.9 provides a comparison of HIPAA and the various other state and federal laws pertaining to privacy and security.

Other public sector initiatives pertaining to healthcare information privacy and security are considered in Section 71.6. Descriptions of the government-instigated HITECH push to move the United States toward so-called "meaningful use" of a new healthcare model that relies on secure, nation-wide, electronic sharing of protected healthcare information are given in Section 71.6.1. Section 71.6.2 provides an overview of the several other regulations and guidelines—other than HIPAA and HITECH—that can apply to healthcare information protection throughout the U.S. Defense Department (DoD). A glance at yet other government sector initiatives pertaining to healthcare information protection is given in Section 71.6.3.

Section 71.7 focuses on private sector initiatives pertaining to healthcare information privacy and security. The Health Information Trust Alliance (HITRUST) industry initiative that creates industry standards and compliance evaluations for the healthcare sector is described in Section 71.7.1, while industry efforts associated with medical device security are described in Section 71.7.2.

Finally, as technology evolves, healthcare security and privacy considerations and policies are impacted. Section 71.8 examines the healthcare information protection aspects of several emerging technologies that are likely to be introduced into the healthcare sector.

71.2 MOTIVATION FOR HEALTHCARE INFORMATION PRIVACY AND SECURITY.

There are many motivations for providing good security and privacy for the various types of transactions and information that flow through, or are stored, processed, and/or displayed by, different hardware, software, or virtual and human entities and services that are either within the healthcare sector or outsourced to parties outside the healthcare sector.

One motivation for a healthcare entity or individual provider to comply with applicable laws and regulations is to reduce the risks of financial and/or criminal penalties including regulatory fines and litigation fees. Such penalties may be incurred from (a) exposure of any patient's personal, private information, or (b) compromise of healthcare-related business transactions and business data. Government imposed penalties for either case can be extensive and provide motivation for improving healthcare privacy and security. Private litigations, especially class action privacy breach suites, can be even more costly. Summaries of total costs and penalties are examined in Section 71.3.2 (including Exhibit 71.1) and Section 71.5.6.4.

Consequences to legitimate patients can be extensive, ranging from financial liability for services fraudulently received by thieves, to higher insurance payments to cover stolen care services, to denials of care due to fraudulent attainment of care service limits. Patients want the privacy and security to prevent such financial tolls.

Another motivation is to reduce the risk of a provider, pharmacist, or carrier making potentially life-threatening (and litigable) medical diagnoses or treatment plans, prescription alterations, or service denials. Altered or missing medical information can lead to endangering patients via lower quality of care resulting from bad decision making and potentially devastating consequences.

Another motivation for a healthcare entity is to reduce the capital and labor costs associated with items such as additional security reviews, consulting fees, additional security awareness training, altered human processes, altered facility processes, and/or IT repairs necessitated by healthcare breaches. An associated motivation is to reduce the costs of providing timely notifications and possibly credit monitoring services to victims of information exposure.

While patients have not yet been studied as potentially fickle healthcare customers, a motivation to provide adequate privacy and security in other fields is to avoid the loss of customers when an information breach, or history of breaches, tarnishes a company's image and reputation.

The above motivators result from the existence of imposed negative consequences arising from exposing personal data. On the other hand, there are healthcare entities that are motivated by maintaining or enhancing their reputations as parties that provide superlative protection to healthcare information and transactions.

An exemplar sampling of the kinds of attacks, threats, and vulnerabilities that can provide such motives are shown in the exhibit in Section 71.3.

71.3 HEALTHCARE SECTOR THREATS AND VULNERABILITIES LANDSCAPE.

As the healthcare field embraces new networking and information handling technologies, vast sets of threats and vulnerabilities to health-related systems and information pose greater and greater risk to providing and administering patient care. A synopsis of the growing occurrence of healthcare data breaches appears in Section 71.0. As examined in Section 71.0, a wide spectrum of types of attacks are used to breach healthcare information and lead to several types and sizes of consequences. The attack landscape is constantly evolving and expanding (Section 71.0), keeping the healthcare sector in a nondiminishing state of risk.

EXHIBIT 71.1 Healthcare-Related Breaches

Responsible Organization and Partners	Date	Exposure Type	Exposed Data	Number of Records	Impact
Computer Sciences Corp (Medicaid contractor), NC & VA	3/1/13	Contractor employee lost USB drive with personal information of NC Medicaid providers	Provider names, SSNs, addresses, birth dates; no patient information	50,000	NC HHS directs CSC to conduct outside review of its security, provide victim notification and hotline for provider inquiries.
Lucile Packard Children's Hospital (LPCH), Stanford University, Palo Alto CA	1/13/13	Theft stole password-protected, unencrypted laptop from physician's car parked off-campus	Patient names, birth dates, med record numbers, med descriptors, clinical data from 2009 patient care and research	57,000 (approximate)	LPCH established investigation and strengthened data protection policies, controls, and encryption. Previous 2010 LPCH laptop theft resulted in $250 K fine (reduced to $1 K) issued by CA Department of Public Health after HHS investigation. Costs of investigations, patient-optional identity protection services, & call center for patient inquiries. Unknown potential future fine by CA Department of Public Health.
Utah Department of Health (UDOH)	1/13/13	Contractor employee loses unencrypted portable USB memory stick	Medicaid enrollee names, ages, prescriptions	6,000 (approximate)	Unknown. Potential impacts could be to improve employee awareness of security policies and/or to deploy technology to prevent copying of unencrypted sensitive files.
U.S. Air Force Health Longitudinal Technology Application clinic data repository	1/12/13	Bad software upgrade	Military personnel health records in 412 medical clinics and 65 military hospitals	9,000,000 (approximate)	Global health records system outage; service degraded; only local records available

(continued)

EXHIBIT 71.1 (*Continued*)

Responsible Organization and Partners	Date	Exposure Type	Exposed Data	Number of Records	Impact
Utah Department of Health (UDOH)	4/2/12	Thieves gained access to server with misconfigured authorization component and download healthcare client personal data	Medicaid claims & Children's Health Insurance Plan records, including patient names, SSNs, addresses, birth dates, provider names, & tax ID numbers	780,000	HHS investigation ongoing. UDOH reputation damaged, as clients feel violated by personal information loss; UDOH enhanced net monitoring & intrusion detection. Utah analyzing protection of all servers and IT policies & procedures. Public trust in emerging Utah statewide health information exchange needs rebuilding. Costs of victim mail notices + 1 year of free credit monitoring services. Utah Department of Technology Services head fired.
Emory Healthcare, Atlanta, GA	2/7/12	Emory lost 10 backup tapes containing data from an obsolete system deactivated in 2007	The backup contained 17 years of surgery patient data including names, SSNs, surgery dates, diagnoses, procedure codes, surgeon & anesthesiologist names, devices used in surgery	315,000	Emory paid costs of victim notification via postal service plus one year of credit monitoring and identity protection services. Emory reinforced, clarified security and privacy policies & procedures; conducted appropriate inventories.
St. Jute Medical Center, Mission Hospital, Queen of the Valley Medical Center, Santa Rosa Memorial Hospital, Petaluma Valley Hospital	Early 2011 through Feb 2012	Incorrect security settings allowed external Google and Yahoo! searches	Patient names, demographic data, medical info, (e.g., blood pressures, lab results, med allergies, BMI, etc.)	31,800	Cost of patient mail notification.

Howard University Hospital, Washington DC	1/25/12	Thief stole contractor's personal, password-protected laptop from car	Patient info including SSNs, addresses, ID numbers, medical record numbers, birthdates, admission & discharge dates, diagnosis medical codes, medical procedures & patient condition	66,601	Hospital strengthened policies regarding data & laptop encryption, usage and down loading. Costs of patient mail notices and identity theft alert coverage for 1 year.
Indiana University Health Goshen Hospital, South Bend ID	December 2011	Malware used to steal patient and job applicant data	Outpatient & maternity patient SSN, insurance & med service data; applicant name, address, SSN	500 patients; 12000 job applicants	Cost of patient mail notices; hospital provides one year of credit monitoring.
University of Miami, Miller School of Medicine, Miami FL	11/24/11	USB flash drive stolen from car	6 years of patient age, gender, diagnosis, treatment data	1,219	Cost of patient mail notices; theft reported locally and to HHS.
Sutter Health Network billing & managed care system, Sacramento, CA	10/1/11	Computer stolen with unencrypted data	6–16 years of patient names, addresses, birth dates, phone, email, medical records, service dates, diagnoses & procedures, insurance plan name	3.3M patients of Sutter Physician Services; 940 K patients of Sutter Medical Foundation	11 class-action lawsuits filed; rolled into one multi-billion-dollar suit. Voluntarily, Sutter expedites plans to encrypt all desktops while awaiting lawsuit outcome and likely associated monetary damage.
TRICARE Management Activity (TMA) Military health benefits provider for several military clinics & hospitals. Also named was the business partner, Science Application International Corporation (SAIC)	9/13/11	Unencrypted backup tapes stolen from the vehicle of an SAIC employee	The loss potentially exposed 19 years of patient SSNs, addresses, telephone numbers, clinical notes, lab test reports, and prescription information	4901,432	Plaintiffs filed a $4.9B class-action lawsuit against TRICARE, the Department of Defense, and Defense Secretary Leon Panetta for violation of Privacy Act of 1974, failure to encrypt, transport of data off government premises; victims experienced fraudulent financial transactions on credit cards and bank accounts.

(continued)

EXHIBIT 71.1 (*Continued*)

Responsible Organization and Partners	Date	Exposure Type	Exposed Data	Number of Records	Impact
Milton Pathology Associates, P.C. and their partner Goldthwait Associates.	7/26/10	Improper disposal of un-redacted, un-shredded medical records disposed into Georgetown, MA town dump	Patient names, SSNs, medical diagnoses, cancer & other pathology test results	11,000	Unknown. Potential impacts could be to conduct awareness training to follow state and federal health care data disposal laws. Jan '13 $140 K MA settlement for HIPAA and MA civil penalties by all 5 parties, attorney fees & contribution to MA healthcare and financial data protection improvement fund.
Alaska Department of Health and Social Services (DHSS)	10/12/09	Thief stole portable USB memory stick	Personal info of state Medicaid beneficiaries	501	DHSS investigation reveals inadequate information security practices that increased size of HHS fine. $1.7M settlement of HHS HIPAA civil fine + costs to institute risk analyses & management, conduct security awareness training, implement adequate IT security controls.
BlueCross BlueShield of Tennessee (BCBST)	10/7/09	Intruder stole 57 unencrypted hard drives from locked closet in training center in a leased facility	The hard drives contained customer support call recordings of BlueCross BlueShield Identifier, name, birth date, Medicare Health Insurance Claim number, SSN	1023,209	Department of Health and Human Services imposed HIPAA Security Rules & HITECH breach notification financial penalties and regulatory correction directives of $1.5 million. Additionally, BlueCross BlueShield of Tennessee spent $17 million on investigations, pen testing, notifications, employee awareness & security training, installation of data encryption, and free credit monitoring for victims.

APA Medical Practice, Atlanta GA	April 2010	Disgruntled former employee steals and deletes patient database	Patient names, telephone numbers, addresses	Unreported	APA database purged impacting provider patient care on case-by-case basis; APA database given to competitor for marketing campaign.
University of California at Los Angeles	2005–2009	Hospital staff snooped on medical records	Undisclosed amount of personal information viewed	2 celebrity patients	HHS imposed civil monetary penalty for HIPAA violation. $865 K penalty settlement and a multi-year program to prevent future violations.
Hospice of North Idaho	June 2010	Theft stole laptop from home hospice employee's car	Unencrypted sensitive personal patient info	441 patients	Hospice undertook several security and privacy improvement activities, e.g., risk analysis, stronger security controls, policies & procedures, laptop encryption, HIPAA training. Hospice undertook several security and privacy improvement activates, e.g., risk analysis, stronger security controls, policies & procedures, laptop encryption, HIPAA training.
Veterans Affairs	May 2012	Transmission of unprotected (i.e., unencrypted) electronic health records	Veterans' electronic health records from medical centers and outpatient clinics	Unknown	Veterans Affairs CIO Roger Baker and undersecretary Robert Petzel signed waivers delaying encryption. No immediate harm to patients known.
Accretive Health, Inc. (business associate)	7/1/11	Laptop stolen from Accretive Health employee	Unencrypted patient information	23,000 (approximate)	Minnesota Attorney General successfully sued Accretive Health, Inc. On July 30, 2012, Accretive Health agreed to pay $2.5 million in damages.

71.3.1 Healthcare Data Breaches are Growing. There is a growing body of evidence from the rise in healthcare information breaches that the need for comprehensive strategies for protecting healthcare-related data is aggressively growing. A snapshot of this accelerating need is presented in this section.

In the healthcare sector, a malware event occurs, on average, about once every four minutes with only modest variability from month to month.[3] Occasionally, healthcare sector attacks are event driven, such as the surge in healthcare sector malware attacks after China identified healthcare as one of its top priorities in its long-term science and technology development plan.[4]

As more medical records become electronic and as new IT storage, computing, and communication technologies appear in healthcare entities, the risk of healthcare breaches of personal health information (PHI) and/or breaches of other types of personally identifiable information (PII) found in a healthcare setting has increased in frequency and magnitude. The number of records stolen in a single breach has increased. Often months pass before a breach is discovered. It appears that the trend in focus of attacks in the healthcare sector is paralleling that seen in other sectors; namely personal data, especially medical identities, are being targeted more, while financial data are of less interest.

It is said "healthcare is one of the most-breached industries"; in the mid-2000s nearly 40 million patient records were breached in a four-year period.[5] But, in 2011 alone, breaches doubled from the previous year.[6] Typically, about three to nine[7] times more healthcare-related breaches occur than financial industry breaches. About 19 million patients' health records were breached via nearly 400 breach incidents. About 30 percent of healthcare entities experienced a data breach in 2012, with more than two-thirds experiencing more than one breach.[8] Furthermore, it is estimated that nearly all small healthcare offices have been breached. Medical identity theft occurred in nearly a third of these small office breaches.[9]

Motivations for initiating medical record breaches range from attackers seeking monetary gain for new employers[10] to attackers selling SSNs, medical IDs, and medical records. Indeed, in 2012 stolen medical records that include a health insurance ID number fetched $50[1] each on the underground market—an amount greater than the typical $1 value placed on stolen financial records that include a Social Security number. In 2012, each medical ID theft cost more than $20,000[11]—about ten times the cost of a regular identity theft—and took twice as long to discover and to resolve as regular ID theft.[12] All in all, though, it is believed the healthcare sector does not yet have a good picture of the major threat and attack motivators.[13]

The U.S. Department of Health and Human Services (HHS) maintains an extensive, formal list and synopses of healthcare information breaches.[14] Breaches affecting 500 or more people must now be reported to HHS. In practice, HHS postings appear to list many of the smaller healthcare-related data heists. The heists reported may not often appear in HHS's list in a timely fashion.

71.3.2 Scope of Attack Vectors and Breach Consequences Is Large. The numbers of healthcare information breaches and healthcare service identity thefts are enormous. Recently, the federal government investigated over 27,000 of the 77,000 complaints of healthcare information breaches it received; it monitored for misuse nearly 300,000 Medicare beneficiary IDs that were compromised.[15]

Exhibit 71.1 reveals examples of the diversity of types and sizes of data heists or exposures occurring in the healthcare sector. The exhibit includes many larger breaches

not yet reflected in the HHS list. Many of the tabulated breaches are associated with some of the highest volumes of healthcare data compromise and the most significant associated remediation costs and penalty fines.

The exhibit shows the custodial source of the compromised data, characteristics of the data that were compromised, and the consequences of the data compromise. The characteristics of the data compromised include the date of the compromise, the type of data compromise that was experienced, the type of data that were exposed, and the number of records that were exposed. The consequences of the data compromise are described in terms of the impact of the compromise and financial penalties associated with the data compromise.

The exhibit shows that healthcare information breaches know no preferred source of compromised data; compromises occur in academia, private hospitals, military hospitals, small physician practices, large insurance carriers, small carriers, and more.

Inspection of the exhibit reveals the large diversity of threats and attacks that result in PHI loss. The threats and attacks range from social engineering attacks, to insider attacks, to poor patch management in medical device software, to outright theft, to ransom, and more. All too often, unauthorized requesters successfully extract the private, personal medical information of unsuspecting individuals—sometimes as easily as listening into minimally secured wireless networks in conjunction with a drive-by heist—or thieves steal patients' insurance billing information for financial benefit.

Common among the cause of many breaches is the theft or loss of portable devices (e.g., USB drive, laptop, backup media) that contain patient personal information. Missing devices with compromised PHI have been found in city dumps and pawn shops.

One of the underlying trends is that the data exposed by such breaches are often made available for exposure by unsuspecting healthcare entity employees who download PHI into less-protected, inadequately permissioned, poorly secured personal devices. Such downloads are typically not made for malicious reasons, but the data become more vulnerable nonetheless.

The exhibit also shows a frequently recurring commonality of the type of personal information exposed by the data that were compromised. Such common data items include name, SSN, birthday, address, and so on.

Ransoming healthcare information has emerged as a new threat that may be targeted[16] at small and midsize healthcare entities. In the future, as large volumes of medical data are aggregated into centralized, networked, health information exchanges, the potential for the formation of a cybercriminal's treasure trove bonanza increases dramatically.

Inspection of the exhibit also reveals a wide spectrum of consequences associated with PHI compromise. Lower price consequences include, for example, the costs of victim mailings, victim credit monitoring, and victim identity theft surveillance. On the other hand, the total cost of other breaches includes not only such victim-oriented remediation activities but also the costs incurred from federal or state regulators' penalties and fines. The average cost incurred by a custodian of breached health information is currently about $150 per record breached. But enterprising attorneys representing patients whose health information was breached are looking to set precedents for multimillion/multibillion dollar class-action lawsuits against covered entities that violate California's Confidentiality of Medical Information Act, which sets damages for health information exposure at $1,000 per person per privacy violation.[17] Such breaches can lead to very significant aggregate per-incident costs to the affected healthcare entity.

The consequences to a patient arising from a healthcare data breach are typically similar to consequences experienced in other industries. Often, stolen identity data are

used to apply for credit consumed in the name of the victim from whom the identity information was stolen. In the case of stolen medical identity data such as carrier policy numbers associated with specific patients, the stolen data can be used illicitly to receive medical care billed to the person whose medical identity was stolen.

A new type of emerging consequence to a healthcare information breach is that individuals who received data breach notifications are also becoming victims of fraud. It was estimated in 2012 that 25 percent of breach victims become fraud victims, often because cyberthieves use stolen victims' Social Security numbers to defraud the victims' financial institutions. In the Utah Department of Health data breach, the $2 to $10 million being spent to recover from the breach is expected to be dwarfed by the more than $400 million lost in the expected 122,000 cases of fraud, where each case is expected to result in an average loss over $3,300.[18]

Reputation tarnish for a care provider that is a source for compromised PHI is a consequence that is hard to quantify; but consumers of services from the Oregon Health & Science University are likely to seek care elsewhere now that this covered entity has been involved with three significant breaches in four years.

The consequences of a PHI breach on a health entity's corporate executives can be significant. While C-level imprisonment is possible under Federal regulations specified in the Health Insurance Portability and Accountability Act (HIPAA) examined in detailed throughout this chapter, the consequences of healthcare data breaches on the sources of compromised data more typically range from expensive civil law suits, fines, and penalties to expensive changes in security and privacy policies and procedures.

The exhibit demonstrates that Federal authorities and state authorities are serious about enforcement of penalties for losing PHI. Indeed, HHS and other government agencies are providing precedents that government bodies will go after cases of large or small volumes of data compromised, as well as large or small violators, with equal vigor.

One of the first punitive actions of the HSS Office for Civil Rights (OCR) is shown via the BlueCross BlueShield of Tennessee (BCBST) entry in Exhibit 71.1. OCR applied the enforcement tool of the HIPAA Breach Notification Rule (described in Section 71.7). OCR used the precedent-setting BCBST medical record data breach penalty as a significant example to demonstrate to the healthcare community the very serious consequences of not having a high-priority, carefully designed, up-to-date, and monitored HIPAA compliance program.[19]

On the other hand, HHS also made an example of a small hospice that compromised personal information from only a very small number of patients.

The exhibit also reveals that in the case of Accretive Health, Inc., the state of Minnesota has established precedent that governments will undertake compliance enforcement actions with significant penalties against HIPAA-designated Business Associates of covered entities.

71.3.3 More Attack Vectors are Emerging. There are other attack vectors on health information and healthcare entities not shown in the table of examples. For example, thieves used fraudulent Automated Clearing House transactions though duped third parties to steal $1 million from a Washington State hospital.[20] Indeed, independent of applicability just to the healthcare sector, there is a constantly evolving and expanding threat landscape. As a consequence, new security approaches will likely be needed to respond effectively to the changing threat picture.

Furthermore, there will always be some new wave of vulnerabilities; the current wave of vulnerabilities associated with IT trends regarding mobility, social media,

cloud computing, consumerization, and so on will undoubtedly be replaced by newly emerging vulnerabilities from other emerging IT trends.

For example, evidence is growing that in vivo electronic medical devices that communicate with external devices can become vulnerable to computer and/or process control attacks.[21] Researchers have shown that wearable and implantable medical devices such as insulin pumps, pacemakers, defibrillators, and more can be wirelessly hacked and device settings can be changed—even to the point whereby a rogue-controlled medical device can be configured to deliver fatal consequences to the associated patient. There is concern that constraints on in vivo power may limit the ability of an implanted or wearable medical device to provide the necessary security controls, such as encryption and authentication, to protect communications to/from the human body. Of further concern, the Department of Veteran Affairs within the U.S. Department of Defense (DoD) reported that hundreds of medical devices were infected with malware over a two-year period.

Other unexpected vulnerabilities that have the potential to be exploited include attacks on a healthcare entity's facility control systems, for example, those used in conjunction with electronic locks, surveillance cameras, lighting, alarms, hospital kitchen refrigeration, and so on.[22]

Drive-by thieves can eavesdrop on patient-specific information if such information is conveyed via an unprotected wireless network at a care provider's facility or home or a patient's home, or is displayed on an unattended computer screen. A lucrative primary target of such a healthcare-specific attack is the theft of patient-specific billing information. An unscrupulous care provider billing for unrendered services can use a patient's insurance ID number to perpetrate fraud. Alternatively, a patient without insurance can use a stolen patient ID number to obtain free, or nearly free, healthcare. Such fraudulent activity impacts healthcare premium costs, patient benefits, and patient medical histories needed for accurate, future diagnoses.

As patient medical records are made electronic and are aggregated into large data stores or healthcare information exchanges, the potential escalates for massive medical records losses to occur in one fell swoop. Challenges in maintaining care providers' patient care continuity and medical business continuity in response to potential cyber Pearl Harbors will take center stage.

Additionally, there is evidence that elements of the U.S. healthcare sector and the National Health Information Network being used to prototype the exchange of electronic medical records are under potentially coordinated attack by multiple government elements of a rogue nation state.[23] Primary attacks appear to include cyberspying on U.S. healthcare-related technology intellectual property and cyberespionage on healthcare business processes used by healthcare providers and insurance carriers. At risk are billions of dollars effectively being stolen, for example, via espionage and theft of pharmaceutical, medical device, and medical/biotechnology R&D and drug trials,[24] or via illegal pilfering and use of medical insurance. Furthermore, the potential exists for cyberadversaries to use such cyberpilfered knowledge to cause significant harm to patients—including loss of lives. Even more troubling is that there are emerging attacks, currently on other industries, where the motive for attack is solely to cause destruction.

71.3.4 The Bottom Line. Recent analyses[25] show that healthcare data breaches at the time of the writing of this chapter were, at least temporarily, diminishing in numbers of breach incidents and in the severity of the breaches in terms of the numbers of patient PHI records exposed. While providing an aurora of encouragement that perhaps

the government regulations and imposed penalties described later in this chapter were starting to make improvements in protecting PHI, it is too early to conclude a corner has been turned in healthcare data protection or whether there was just a statistical calm before the next storm of exposures.

Whether there is or is not a long-term lull in healthcare data exposure, complacency about healthcare data protection is not an option. The prudent conservative approach is to assume that the historical epidemic of security will resume with a vengeance, that every healthcare entity is a target, and that the healthcare entity of concern to you will be the next target. Those charged with ensuring the security and privacy of healthcare-related information must be constantly vigilant about new threats and vulnerabilities that emerge in order to assess whether fielded security protections are adequate to withstand newly emergent threats and vulnerabilities.

Many organizations provide listings of recently discovered, general threats not necessarily specific to the healthcare sector; one,[26] for example, also enumerates costs incurred in recovery and information that can help in reducing risks. Furthermore, NIST maintains a database of discovered vulnerabilities (sometimes applicable to technology used in the healthcare sector) that can help guide vulnerability remediation efforts.

Perhaps most troubling is the emerging thought that privacy could be made impossible for certain types of medical information because of a new style of so-called "DNA hack."[27]

71.4 HEALTHCARE INFORMATION PROTECTION REQUIREMENTS IN U.S. LAWS.

This section contains brief descriptions of several of the state and federal laws in the United States that stipulate privacy and security requirements pertinent to patient care and healthcare sector business processes.

Privacy requirements were initially established for personal information maintained by the Federal Government. Privacy requirements were subsequently extended to applicability beyond government-held personal information.

Privacy of PII first became a requirement in the United States via The Privacy Act of 1974. This law applies only to those databases established, used, and maintained by any agency of the Federal Government. As such, this seminal law established privacy requirements pertaining to personally identifiable health information collected by any agency of the U.S. Government, such as Medicare treatment records entrusted to the HHS Centers for Medicare and Medicaid Services.

This law established the precedent that disclosure of PII (in this case by a federal agency) requires written individual consent of the person whose PII is to be disclosed. The law also specifies the few specific exception situations wherein disclosure consent is not required. The law specifies other privacy requirements and rights applicable to PII in databases wherein a Federal Government agency is the custodian of the database. The law also specifies the interworking with the Freedom of Information Act. The precedents established by the Privacy Act have since been incorporated into all subsequent public or private privacy rules applicable to a variety of industries, including healthcare.

Since the time that the Privacy Act was established, several presidential administrations have issued plans and policies pertaining to healthcare information entrusted to the Federal Government. For example,

- **Presidential administrations** issued policies, plans, and positions wherein the healthcare sector is deemed a critical national infrastructure. It needs to be protected. It needs to continue operating because the nation and economy depend on it.

- **The Health Care Finance Agency (HCFA)** established policy for incorporating security functionality and for appropriately using the Internet to convey healthcare information protected by the Privacy Act as well as other sensitive HCFA information.
- **The Office of Management and Budget** promulgated policy recommending all federal agencies, including those operating as healthcare entities for which HIPAA is also applicable, to implement protection mechanisms, including encrypting all data stored in mobile devices and establishing database access logging.
- **The Federal Acquisitions Regulation Council** issued a rule specifying that procurement officers at all federal agencies are required to incorporate IT security requirements in their IT acquisition planning. This policy applies to federal agencies that act as entities covered under HIPAA.
- **The Executive Branch** issued the Homeland Security Presidential Directive/HSPD-12 requiring government employees and contractors working at government agencies—including those operating as entities covered under HIPAA regulations—to use specific government IDs issued by trusted authorities only to identity-proven individuals.

There are yet other guidance and law that pertain to federal/DoD healthcare facilities and the healthcare information generated and used therein. These other healthcare information protection rules for government-entrusted personal healthcare data are presented in Section 71.6.2.

Some other laws are not specifically concerned with security, but compliance with such laws may infer the need for best-practice security implementations addressing implied security requirements associated with certain types of healthcare-related business information. For example, under the False Claims Act, healthcare providers submitting fraudulent claims can be punished with triple damages, penalties of $5,000 to $10,000 per wrongful billing, and exclusion from government insurance programs. Without implementation of adequate information security, electronic billing of health insurance carriers by healthcare providers can open doors to malicious attacks on healthcare claims so that they appear to be wrongful, prosecutable billing claims.

Currently, several other laws, standards, regulations, policies, and guidelines exist to protect the privacy of PII that is not maintained by the Federal Government. Such PII includes the specific PII used in healthcare.

Many of these rules and suggestions apply to protecting business processes, including the business processes associated with the delivery of healthcare. Some of the better-known guidance includes the Payment Card Industry (PCI) Data Security Standard (DSS) 2.0 and the Gramm-Leach-Bliley Act's (GLBA) Personal Financial Information (PFI).

Many of these overlap somewhat, giving the practitioner a point of reference when addressing healthcare data protection. HIPAA contains 18 distinct data elements. Some elements map closely to PCI and GLBA, in that HIPAA includes *name* as an element, whereas both PCI and GLBA are more specific. PCI and GLBA require first name or initial in conjunction with last name. Furthermore, PCI and GLBA must accompany a credit card number or financial account number respectively.

This minor difference broadens the scope of HIPAA such that either first name or last name could constitute *name* in the HIPAA context. This broadening extends to other, less standard, elements such as vehicle identification number (VIN) and telephone number. The key to understanding HIPAA's reach is to compare its purpose with PCI,

GLBA, or state data privacy laws like Massachusetts 201 CMR 17.00. HIPAA seeks to keep information about the patient private and knowledge of health data in and of itself is the act to be prevented. Others seek to prevent a subsequent action after an actor obtains vital combination data.

The HIPAA law and its various subsequent rules extend PII protection to personal healthcare information (PHI) created and used not necessarily just for healthcare business processes in both the private and public sectors but also for private sector and public sector patient care processes. The security, privacy, and breach notification requirements stipulated in HIPAA are examined in detail in the next section.

Additionally, recent updates and new state breach notification laws and state data protection laws govern much of the data considered PHI. Nearly all U.S. states and territories now have their own specific laws pertaining to breaches of PII.[28] Such laws typically specify requirements on notifications that must be provided when information breaches are detected, what specific types of PII disclosure require such notifications, and under what conditions (e.g., exposure of encrypted data) breach notifications are not required. For example, California's Security Breach Information Act (Senate Bill 1386) requires any company—including those that are entities covered under HIPAA—to disclose potential security breaches of specified types of PII such as a person's name, specified information about a person's various financial accounts, a person's Social Security number, specified information about a person's health insurance, and a person's driver's license and state ID number.

While many of these state laws maybe intended to be used when "consumer" PII is compromised, it is not far-fetched to expect that law enforcers may interpret patients as consumers (of healthcare) with which exposable healthcare-related PII is associated.

A comparison of portions of the laws identified in this section with the HIPAA law is given in Section 71.5.9.

71.5 THE HIPAA AND HITECH HEALTHCARE INFORMATION PROTECTION REGULATIONS.

The Health Insurance Portability and Accountability Act (HIPAA) of 1996[29] is a five-part U.S. federal law that pertains to many aspects of healthcare and its administration. HIPAA goals include (a) improving portability and continuity of healthcare insurance when a patient changes jobs, (b) minimizing waste, fraud, and abuse in healthcare insurance and healthcare delivery, (c) ensuring individual patient information privacy through security means, and (d) simplifying healthcare insurance administration. The Administrative Simplification (Title II, Subtitle F part of HIPAA) looked to improve the efficacy and reliability of the healthcare system and to reduce its costs.

The initial focus of HHS with regard to Administrative Simplification was to establish key pieces of an infrastructure for electronically exchanging certain healthcare-related financial and administrative transactions. The national standards[30] that were established included use of standard, specifically formatted electronic data interchange (EDI) transaction types with specific medical code sets for transaction contents, unique health identifiers for employers and care providers, health information security, and electronic signature. This portion of Administrative Simplification is administered and enforced by the Centers for Medicare and Medicaid Services.

But of interest to this chapter, Title II of HIPAA is one of several laws, standards, regulations, and guidelines that exist for protecting the healthcare information associated with patient-specific information. HIPAA in combination with the healthcare

information protection aspects of the Health Information Technology for Economic and Clinical Health Act (HITECH) of 2009 provide federal standards and regulations, as well as legal underpinnings, for the provision of privacy and security protection to personal health information (PHI) and to medical identity.

HIPAA instigated moving the healthcare industry toward widespread use of information technology; personally identifiable electronic health information whose privacy was ensured by security technology and methods; the sharing of such information among authorized healthcare parties; and a single, common, secure electronic transaction environment for electronic billing.

To affect the intents of Administrative Simplification with regard to the topics of interest to this chapter, HHS developed several so-called HIPAA Rules. These HIPAA Rules are the regulations containing requirements that pertain to protecting PHI and to notifying individuals if their PHI was compromised to unauthorized parties.

The privacy and security regulations for protecting PHI apply to PHI while it is (a) in the custody of so-called covered entities and their business associates, (b) in transit between covered entities, and (c) in transit from covered entities to others such as business partners. Security and privacy are related in that in general, privacy depends on security: Appropriate security prevents theft or leakage of private data.

Regulations include developing, documenting, and implementing privacy, security, and breach notification policies and procedures (P&P) that meet HIPAA-stipulated requirements in these three areas, implementing technology and methods for addressing the specified HIPAA requirements, training staff about HIPAA requirements, guaranteeing patient rights regarding their PHI, and appointing privacy and security directors to monitor and to enforce privacy and security matters.

Summaries of the three specific HIPAA Rules appear below. First, the guiding principles for addressing healthcare information protection are given in Section 71.5.1. Descriptions of the original HIPPA Privacy Rule, the HIPAA Security Rule, and the HIPAA Breach Notification Rule are given in Sections 71.5.2, 5.3, and 71.5.4, respectively.[31] Regulations applying to the enforcement of the above three HIPAA Rules are given in Section 71.5.5.

These original HIPAA Rules were subsequently updated under the authority of the HITECH Act. The updates and modifications to the HIPAA Privacy Rule, the HIPAA Security Rule, the HIPAA Breach Notification Rule, and the enforcement of HIPAA Rules are described in Section 71.5.6.

Observations associated with the fielding of the HIPAA Rules and demonstrating compliance of various healthcare entities to these Rules appear in Section 71.5.7.

Various government-issued guidelines relevant to HIPAA and healthcare information protection exist. Examples are given in Section 71.5.8.

States within the United States can establish laws and regulations that are more, or less, comprehensive or restrictive than the federal HIPAA and HITECH laws and regulations. Examples of state-specific privacy and security laws are compared to HIPAA in Section 71.5.9.

Still other federal standards and guidelines exist that can impact the healthcare information security and privacy aspects of military hospitals (see Section 71.7.3). Standards developed by industry can impact the healthcare information security and privacy aspects of certain healthcare parties (see Section 71.7.4.1) as well as medical device manufacturing (see Section 71.7.4.2).

Finally, evolving and potentially emerging new government initiatives pertaining to security and privacy for healthcare information are identified and briefly examined in Section 71.7.8.

71.5.1 HIPAA's Privacy and Security Fundamental Principles. The HIPAA approach for healthcare information protection focuses on stipulating privacy and security requirements rather than specifying privacy and security solutions.

The privacy and security requirements apply to (a) "covered entities" and (b) to all individually identifiable, patient health information handled by such covered entities.

Covered entities include:

1. Health insurance carriers (sometimes referred to as "health plans," including private entities as well as government-funded entities such as the Centers for Medicare and Medicaid Service that act as insurance carriers);
2. Healthcare providers (including, e.g., institutions like medical centers, private physicians, dentists, pharmacies, and government agencies that provide healthcare services such as the Federal Employee Health Benefits Program, and DoD's TriCare Program), only if such healthcare providers conduct certain specified electronic transactions;
3. Healthcare billing clearinghouses (entities between providers and payers that can, if necessary, translate covered entities' transactions between HIPAA standard formats and nonstandard/proprietary, payer-specific or provider-specific formats);
4. Business associates, contractors, and service providers of covered entities (including government contractors such as those contracted for processing Medicare claims).

More detailed examples of which types of organizations are considered to be "covered entities" are given by HHS (www.hhs.gov/ocr/privacy/hipaa/understanding/coveredentities). Not every organization that generates, processes, or holds personally identifiable health information is a "covered entity" in the eyes of HIPAA. Tools to determine whether any specific entity is a "covered entity" exist.[32]

The privacy requirements stipulate how PHI is controlled. Legally authorized uses and disclosures are defined. Privacy requirements pertain to PHI created, stored, or conveyed in any form. For PHI in electronic form, privacy protection largely depends on HIPAA-stipulated security requirements.

The security requirements apply only to PHI in electronic, "computerized" form.[33] By securing electronic transactions that convey PHI between healthcare IT systems and by securing storage and processing of electronic PHI, abuse and fraud can be minimized and privacy of such information can be increased. The security requirements can be met by common, generally accepted, widely used security principles and practices.

HIPAA does not specify a one-size-fits-all privacy and security solution, applicable to all situations. Although every covered entity that electronically receives, maintains, processes, or transmits PHI must comply with applicable HIPAA privacy and security requirements, HIPAA does not specify how to meet these requirements. Privacy and security requirements may be met by covered entities by appropriate combinations of P&P, measures, security technology, and other ways deemed appropriate by each covered entity.

By focusing on privacy and security *requirements* rather than *solutions*, HIPAA offers (a) flexibility for differing covered entities to employ differing security solutions, (b) scalability of differing security solutions to work in covered entities of differing sizes, and (c) technology neutrality so that covered entities may use the newest applicable security solutions.

As such, HIPAA allows different covered entities to meet HIPAA information protection requirements by whatever strategy is reasonable and appropriate to each covered entity at any given point in time. Particular solutions may vary from situation to situation, but each solution must meet the HIPAA requirements. Some healthcare enterprises may need to implement more sophisticated information protection solutions than others. Small, less complex healthcare provider offices may have more limited privacy and security P&P and training requirements than those of large, more complex medical centers. Healthcare enterprises may need to change aspects of their information protection strategy as they change in size.

Covered entities must each assess their own potential risks to the PHI they handle. They must develop, document, implement, and maintain appropriate information protection P&P and measures to address their unique risks and situations. They can use risk-management techniques to help decide what strategy of potentially overlapping safeguards is reasonable and appropriate to their situations. All analyses, decisions, and rationale for decisions must be documented. HIPAA allows covered entities to select industry-acceptable, best-practice information protection strategies. Covered entities need periodic assessments to ensure continuing appropriateness and adequacy of their solutions to meet HIPAA requirements.

71.5.2 HIPAA Privacy Rule. HIPAA privacy regulations stipulate requirements for the use, disclosure, and dissemination of personally identifiable healthcare information as well as patient rights in such matters. A synopsis of the HIPAA Privacy Rule and its early corrections and modifications is given in this section. The more recent HITECH updates to the privacy rule appear in Section 71.5.6.2.

The HIPAA Privacy Rule[34] together with early modifications[35] establish a national floor of requirements pertaining to safeguarding privacy of PHI in any form—paper, oral, or electronic—from inappropriate or unauthorized disclosure or misuse. The Privacy Rule regulations apply to PHI created, stored, processed, used, transmitted, or received by covered entities. The privacy regulations limit sharing of PHI by covered entities. They also establish requirements pertaining to rights of patients with regard to understanding and controlling how their healthcare information is used, disclosed, and shared.

The purpose of HIPAA privacy requirements is to ensure patients' personal information is adequately protected and kept private without unduly disrupting the flow of healthcare information needed to provide, to oversee, and to pay for care, to assist healthcare research, and to improve the quality of the healthcare system. These requirements strike a balance between providing patients with greater peace of mind versus the public responsibility to support national medical priorities. Such priorities include protecting public health, identifying public health trends, conducting medical research, improving the quality of care available for the nation, and fighting healthcare fraud and abuse.

Given the massive extent, diversity, and complexity of the healthcare sector, the privacy requirements are comparably comprehensive—and appropriately flexible—to cover the plethora of uses of healthcare information by various covered entities and others outside the healthcare sector. Only a brief introduction and summary of the privacy requirements is given here.

The privacy regulations provide patients with rights to inspect and to correct their personal healthcare information. The original regulations require written patient consent for each covered entity to use a patient's PHI. Disclosure of PHI without patient consent was allowed for only a few, specified reasons. In the 2002 modification[36] of

the HIPAA Privacy Rule, patients' rights of consent were stripped from the Privacy Rule. Patients can't specify who gets to see their PHI and why. But they can get to see their PHI held by all healthcare parties and a history of what disclosures of their PHI occurred without their authorization.

The regulations require business partners of each covered entity—via contracts or other similar mechanisms—to maintain the privacy of any patient information that is provided to them by the covered entity.

The regulations define the generally accepted set of information to be protected and kept private. It includes patient demographic information, common identifiers, insurance billing and payment information, symptoms, examination and test results, diagnoses, treatment, and future care plans that can be associated with or used to identify a patient, and so on. This information may be oral or recorded in any form or medium, such as paper, magnetic tape, computer drive, and so on. There are no privacy requirements that apply to healthcare information that is de-identified via specified ways.

A sampling of the privacy requirements follows. As readers' needs dictate, the full HIPAA Privacy Rule should be consulted for increasingly detailed descriptions of privacy requirements and exceptions that may apply to such requirements.

- **General use and disclosure requirements.** Other than information that is specifically permitted or under the written consent of the applicable patient, a covered entity cannot use or disclose PHI.
- **Specific permitted use and disclosure requirements.** Covered entities may share certain PHI without a patient's formal signed consent only under certain specified situations, including, for example, among healthcare providers involved in a patient's treatment, with third-party payers (e.g., insurance carriers), with the patient, with parties involved with medical reviews assessing the quality of patient care, with parties involved with assessing the performance of care providers and healthcare operations, and with regulatory and public health agencies, ethical researchers, and auditors. For certain specified reasons, a covered entity may share certain PHI with only informal patient permission not involving a signed, formal consent authorization.
- **Required use and disclosure requirements.** When some other law mandates a covered entity to disclose some specific PHI, the covered entity must disclose the PHI without obtaining authorization from impacted individual. Covered entities are permitted—but not required—to share PHI without a patient's consent or permission for 12 national priority purposes outside a patient's healthcare context. Examples of these purposes include public health activities, law enforcement, legal proceedings, protecting victims of abuse, workers' compensation, and so on. Otherwise, the covered entity must either obtain release authorization from the impacted individual or de-identify the individually identifiable health information before disclosing it.
- **Authorization requirements.** A patient-signed, written consent authorization with certain specified conditions and constraints is required by covered entities for any use or disclosure of PHI that is not for treatment, payment, or healthcare operations or otherwise permitted or required.
- **Minimum necessary data requirements.** Covered entities are required to develop and to use P&P to ensure that no more PHI is used or disclosed than the minimum amount needed for each purpose for which information is used or disclosed.

- **Privacy practices notice requirements.** Covered entities are required to develop privacy practice notices, to post these notices, and to collect individual patient acknowledgments of their receipt. These entities must abide by clearly stated procedures that (a) protect the privacy of patients' information, (b) describe how PHI is used and disclosed, and (c) describe how patients can access their PHI.
- **Patient access and control rights requirements.** Covered entities are required to give individual patients the rights to receive a copy of the entity's privacy practices notice and to specify how the entity should communicate with the patient. Patients have the right to inspect, and to obtain copies of, any disclosures and distributions of healthcare information about themselves. Patients have the right to get an accounting of the disclosures and to amend, correct, and restrict the uses, disclosure, and distribution of healthcare information about themselves.
- **Administrative requirements.** Covered entities are required to develop and to implement appropriate privacy P&P consistent with the HIPAA privacy regulations and tailored to fit their size and complexity, to train and to discipline employees regarding privacy P&P, and to designate a privacy official responsible for the entity's privacy practices and activities.
- **PHI safeguards requirements.** Covered entities are required to have reasonable and appropriate administrative, technical, and physical safeguards to prevent intentional or unintentional use or disclosure of PHI in violation of the HIPAA privacy regulations and to otherwise abide by all the allowed and disallowed uses and disclosures of PHI. The HIPAA security regulations (see Section 71.7.1.3) stipulate such safeguards for electronic PHI.

71.5.3 HIPAA Security Rule.

The HIPAA Security Rule[37] is the regulation that stipulates national requirements for protecting electronic PHI that is created, stored, used, processed, displayed, maintained, or communicated by certain healthcare entities. The Security Rule does not apply to situations where PHI is in paper-based or oral form. A synopsis of the HIPAA Security Rule is given in this section. HITECH updates to the security rule appear in Section 71.5.6.1.

These security regulations establish requirements for safeguarding confidentiality, integrity, and availability of electronic PHI. PHI must be protected against reasonably expected threats, vulnerabilities, and hazards as well as against disclosures and inappropriate use not permitted by HIPAA's privacy regulations. The requirements also establish an organizational culture that facilitates the provisioning and maintenance of security safeguards.

There are linkages between HIPAA privacy requirements and HIPAA security requirements. Some requirements overlap or complement each other. For example, privacy and security regulations both contain awareness-training requirements. They can be combined so that employees of covered entities stay aware of both the importance of protecting PHI and the means by which PHI is protected. In other cases, a specific security requirement may need to be implemented to address a specific privacy requirement.

As outlined below, the security regulations stipulate various categories of administrative, technical, and physical security requirements as well as organizational, P&P, and documentation requirements. They require covered entities to assess risk and to provide for the safeguarding of electronic PHI at risk. For each covered entity, there must be fully auditable procedures for controlling access to PHI and for protecting PHI against compromise and misuse.

71.5.3.1 Administrative Safeguards Requirements. These requirements pertain to administrative P&P and actions targeted to manage (a) selection, deployment, and maintenance of security safeguards to protect electronic PHI and (b) roles of employees in protecting PHI. The requirements categories include:

- **Security management process.** Implement P&P for preventing, detecting, containing, and correcting security violations.
- **Assigned security responsibility.** Identify an individual responsible for developing and implementing security-oriented P&P.
- **Workforce security.** Implement P&P for ensuring appropriate employee access to, or denial of access to, PHI.
- **Information access management.** Implement P&P for authorizing access to PHI consistent with privacy access requirements.
- **Security awareness and training.** Implement a security awareness and training program for all employees.
- **Security incident procedures.** Implement P&P for handling incidents that impact systems associated with PHI.
- **Contingency plan.** Establish and implement P&P for responding to emergencies that damage systems containing PHI.
- **Evaluation.** Periodic assessments that the security P&P address the security requirements.
- **Business associate contracts and other arrangements.** Establish legally binding, chain-of-trust agreements with business partners for protecting PHI conveyed to such partners.

A synopsis of the implementation specifications associated with these administrative safeguards requirements appears in the Security Standards Matrix table in Appendix A of the HIPAA Security Rule.

71.5.3.2 Physical Safeguards Requirements. These requirements pertain to physical measures and associated P&P to protect IT systems, facilities, and applicable equipment from physical hazards, unauthorized physical access, and unauthorized intrusions. The requirements categories include:

- **Facility access controls.** Establish and implement P&P to control physical access to IT systems, related equipment, and the facilities that house them.
- **Workstation use.** Implement P&P regarding physical location and use of computers to access electronic PHI.
- **Workstation security.** Implement physical safeguards to control physical access to computers used to access electronic PHI.
- **Device and media controls.** Implement P&P regarding receipt, reuse, and disposal of hardware, software, and electronic media that handle electronic PHI.

A synopsis of the implementation specifications associated with these physical safeguards requirements appears in the Security Standards Matrix table in Appendix A of the HIPAA Security Rule.

71.5.3.3 Technical Safeguards Requirements. These requirements pertain to technological measures and associated P&P to protect and to control access to electronic PHI. A covered entity can determine which security will need to be implemented. The requirements categories include:

- **Access control.** Implement technical P&P for controlling access by people and software to IT systems that handle electronic PHI. Access must be appropriate to the role or job function of the entity seeking access. Subordinate requirements include (a) unique user identification to make users accountable by tracking their electronic activities, (b) emergency, possibly alternative, access procedures, (c) automatic user logoff to terminate inactive session, and (d) encryption and decryption to prevent view or use of PHI by unauthorized entities.
- **Audit controls.** Implement technical and procedural mechanisms that examine and record activities and accesses to IT systems that handle electronic PHI.
- **Integrity.** Implement P&P to ensure the validity of electronic PHI.
- **Person or entity authentication.** Implement procedures to verify legitimacy of people and software attempting to access electronic PHI.
- **Transmission security.** Implement technical security measures to control unauthorized access to electronic PHI in transit over computer networks. Subordinate requirements include reliance on (a) integrity controls to ensure PHI is not modified without detection, and (b) flexibility to establish compatible encryption/decryption methods between PHI senders and receivers.

A synopsis of the implementation specifications associated with these technical safeguards requirements appears in the Security Standards Matrix table in Appendix A of the HIPAA Security Rule.

71.5.3.4 Organizational Requirements. These requirements pertain to using business associate contracts or other official arrangements to establish trust that third parties associated with, or providing services to, covered entities will maintain HIPAA provisions.

71.5.3.5 Policies and Procedures and Documentation Requirements. These requirements pertain to developing all the requisite P&P (as outlined earlier), to developing and using all required documentation and documentation configuration control measures, and to implementing the P&P to comply with HIPAA security regulations.

71.5.3.6 Implementation Specifications. Each stipulated security requirement summarized above must be met by a defined implementation specification stipulated in the HIPAA security regulations, or, if no implementation specifications exist, by measures documented and implemented by each covered entity. NIST Special Publication 800-66[38] identifies (in Table 4 in Appendix D) implementation specifications and companion security controls (detailed in NIST Special Publication 800-53) accompanying each HIPAA Security Rule requirement.

For HIPAA Security Rule requirements met by HIPAA-stipulated implementation specifications, such implementation specifications are individually identified in the

HIPAA Security Rule (and in NIST SP 800-66) as either being "required" or "addressable." Covered entities must comply with all 13 *required* implementation specifications.

For "addressable" security implementation specifications, covered entities must document how they assessed and decided (a) whether each specific *addressable* implementation specification is reasonable and appropriate, or not; (b) whether it applies, or does not, to their specific environment; and (c) whether the covered entity implemented, or not, each addressable implementation specification (or some equivalent alternative measures).

For example, to meet HIPAA's contingency plan requirement, a covered entity is *required* to establish and to implement data backup, disaster recovery, and emergency mode operations plans; and it must also *address* procedures for testing and revising such plans.

Covered entities have options for how to field security solutions for addressable security implementation requirements that are deemed reasonable and appropriate in their environments. Risk, security, and financial analyses should help decide which security technologies or methods best address the security implementation requirements. Covered entities must document and periodically reassess their analyses and rationale for fielding specific security measures.

71.5.4 HIPAA Breach Notification Interim Rule. The U.S. HHS[39] and the Federal Trade Commission (FTC)[40] have companion, harmonized regulations pertaining to notifying individuals whose individually identified PHI is breached. A synopsis of the original Breach Notification Rule is given in this section. HITECH updates to the breach notification rule appear in Section 71.5.6.3.

The HHS regulations apply only to HIPAA-covered entities and to their business associates.

The FTC regulations apply only to certain Web-based businesses, such as vendors of personal health records, as well as certain other parties not stipulated in HIPAA, that provide individuals the voluntary ability to create, to collect, and to manage online stores of their own health information. The FTC regulations also apply to third-party applications that can be used by individuals voluntarily to interact with such health information stores, to create new individually identifiable content for such data stores (e.g., individually obtained blood pressure readings) or to process or to display existing individually identifiable content. The FTC regulations also specify other entities that fall within the jurisdiction of the FTC breach notice requirements.

The federal breach notice regulations specify what is a breach, under what conditions are breach notifications issued, who is responsible for issuing a breach notification, to whom breach notifications must be sent, timing constraints associated with the generation of breach notifications, what information must be contained in the breach notification, precedents among different breach notification laws, and penalties for not following the regulations. The various factors that compose the breach notification requirements are summarized below.

71.5.4.1 Breach Definition. According to the 2009 interim breach notification regulations, a breach occurs when a compromise of electronic PHI (and also paper-based PHI in the case of HHS) takes place and the breached information poses significant financial or reputational risk or other harm to the individual whose PHI was revealed, acquired, accessed, or used by an unauthorized party.

There are three exceptions to the definition of a breach. Breaches exclude inadvertent PHI disclosure between two authorized individuals in the same healthcare entity.

Breaches also exclude PHI disclosures to unauthorized individuals who are reasonably believed not to be able to retain such disclosed PHI. Also excluded are unintentional accesses, acquisition, or use of PHI if such actions were done in good faith within the scope of authority and such actions do not lead to further PHI use or disclosure.

71.5.4.2 Triggers for Breach Notifications. The required condition for generating a breach notification is when any part of any *unsecure*d PHI within the domain of a PHI data custodian is exposed via a breach. The regulations stipulate that *unsecured* PHI is PHI that is decipherable, readable, and usable by parties that are not authorized to access such PHI. HHS has identified technologies and methodologies that can make PHI indecipherable, unreadable, and unusable.[41]

Use of such technologies and methodologies renders PHI adequately protected so that any breaches thereof do not require breach notifications. As such, for example, if encrypted PHI is stolen as a consequence of a breach, the custodians of the breached, encrypted PHI are not required to notify the individuals whose encrypted PHI was stolen. It is assumed that the parties stealing the breached encrypted PHI will not be able to decipher, read, or use the encrypted PHI.

There are other conditions, generally associated with the above breach exceptions, under which breach notices are not required. For one, any health information that has been de-identified according to methods specified in the HIPAA Privacy Rules does not constitute data for which breaches thereof require breach notification. For another, not all unauthorized accesses of PHI lead to unauthorized PHI acquisition and the need to provide breach notices. For example, inadvertent or unintentional access of PHI that does not lead to reading, using, or disclosing PHI may not constitute a breach. Furthermore, an impermissible disclosure of health information to an unauthorized individual who is not able to retain the disclosed health information may not constitute a breach. Similarly, a lost computer or storage device that houses PHI and is subsequently retrieved—and for which it can shown that PHI files were never opened, exchanged, or changed—also may not constitute a breach.

Custodians of breached PHI must be able to demonstrate that all required breach notifications have been sent or that disclosure of unsecured PHI did not meet the definition of a breach.

71.5.4.3 Issuers and Recipients of Breach Notifications. As to who is responsible for generating a breach notice, the HHS regulations require HIPAA-covered entities to notify individuals (and, under certain conditions, HHS and the media) when specified breach notification conditions are met. The FTC regulations require Web-based businesses under FTC jurisdiction to notify individuals (and FTC under certain conditions) when a breach of individually identifiable health information occurs.

When breaches occur within the domain of business associates of HIPAA-covered entities, the business associates are required to notify the covered entity with which a business agreement is established; the covered entities are then responsible for notifying the specific individuals whose PHI was compromised by the business associate.

Similarly, when breaches occur within the domain of service providers for businesses under the jurisdiction of the FTC regulations, such service providers are required to notify the businesses for which service is being provided; the businesses are then responsible for notifying their customers whose PHI was compromised by the service provider.

Several methods of conveying breach notices are permissible. Under certain conditions, custodians of breached PHI may use means other than direct communications to

convey breach notices to specific individuals. Breach notices being sent to the HHS, or updates to previously submitted breach notices sent to HHS, must be submitted electronically via a specified HHS form.

71.5.4.4 Breach Timing Constraints. Breach notices that are to be sent to affected individuals and to the media must be sent as fast as is reasonably expected, but no later than 60 days after breach discovery.[42] Breach notifications issued by HHS business associates or by service providers to FTC-designated businesses must also abide by the as-soon-as-possible, -but-no-more-than-60-days, timing requirement.

The acceptable timing constraints of breach notices sent to HHS or the FTC vary with the number of individuals impacted by the breach. Breaches impacting more than 500 individuals must have breach notices sent within 60 days of breach discovery. Breaches impacting less than 500 individuals may be more leisurely treated in terms of notifying HHS or the FTC.

71.5.4.5 Breach Contents. The contents of breach notices to individuals, to the HHS or FTC and to the media serving the state or jurisdiction where the breach occurred must include: (a) a description of the breach event, including date of the breach occurrence and date of the discovery of the breach; (b) a synopsis of what unsecured PHI was compromised; (c) a statement of what impacted individuals should do to protect themselves from potential harm resulting from the PHI compromise; (d) a description of what the custodian of the breached PHI is doing to investigate the breach, to lessen potential losses that may result from the breach, and to bolster protection against future breach attempts; and (e) contact information for the custodian of the breached PHI.

Breach notices emanating from HHS-stipulated business associates or FTC-stipulated service providers must include a tabulation of the identity of all the individuals whose individually identifiable information was, or is reasonably believed to be, acquired during the health information compromise. These breach notices must also convey the information that is to be relayed to impacted individuals.

71.5.4.6 Breach Law Precedents. For entities that may be subject to both HHS and FTC breach notification rules, compliance with specific HHS requirements is sufficient to claim compliance with corresponding FTC requirements. Establishing such an order of precedent lowers the probability of duplicate breach notices being issued to the same individual.

There is also an order of precedence stipulated between breach notice requirements specified in both federal regulations and state regulations; federal requirements supersede any contrary state regulations. But federal requirements are additive to state regulations that do not conflict with federal regulations. Also, federal requirements do not supersede more stringent state requirements.

71.5.4.7 Breach Notification Policies, Procedures, and Penalties. Entities under the jurisdiction of HHS and FTC breach notification regulations must establish breach notification policies and procedures. Such policies and procedures should include educating employees as to what constitutes a breach, how to document and to analyze potential breaches of unsecured PHI, how to report actual breaches, the intricacies of sending breach notices, and what to do with PHI that employees may have acquired via inadvertent unauthorized access. Such entities must conduct policy

and procedures training of employees, and must appropriately penalize noncompliant employees.

71.5.5 HIPAA Protection Rules Enforcement, Penalties, and Liabilities. HIPAA depends on a compliance-driven process to enforce the regulations established in the HIPAA Privacy, Security, and Breach Notification Rules. Victims of PHI disclosure and individuals believing there has been a violation of any of the HIPAA Rules can start the enforcement process by submitting a PHI disclosure or violation complaint to the HHS Office for Civil Rights (OCR). The required content of, and the rules governing, the complaint are specified.[43] Only certain types of complaints[44] will be considered for enforcement action.

The HHS OCR investigates complaints. It also may unilaterally initiate an HIPAA Rules compliance review of any covered entity. If it is determined that a potential criminal violation occurred, OCR can request the Department of Justice investigate the claim and impose any applicable penalties. The OCR can impose civil (monetary) penalties and the Department of Justice can impose criminal (prison) penalties.

The procedures performed during an OCR complaint investigation are stipulated.[45] OCR decides about the interpretation, implementation, and enforcement of HIPAA regulations. When it appears the covered entity that was the subject of the complaint was not in compliance with HIPAA Rules, OCR often attempts to seek resolution by letting the implicated covered entity take corrective actions to establish compliance.

If the actions are not satisfactory to resolve the disclosure compliant or noncompliance situation, OCR can impose civil sanctions. Covered entities can be penalized $100 per violation, up to $25,000 per year per HIPAA requirement not met.

However, there are large criminal penalties for those who wrongfully handle confidential PHI. Any person, who knowingly obtains, uses, causes to use, intends malicious harm, tries to sell or to transfer or to gain from, or discloses, patient information inappropriately, faces penalties. Penalties range from $50,000 to $250,000 and up to 10 years' imprisonment, depending on seriousness of the offense.

Some argue the real dangers of HIPAA enforcement may arise from loss of patient confidence in the implicated covered entity and possible litigation under contract law. When disclosure of personal information causes damage and the legal system is engaged, HIPAA penalties—especially the criminal penalties—can mean the difference between business failure and business survival for implicated entities. PHI disclosures, whether wrongful or unintentional, have potentials to be devastating. Some argue potential consequences of HIPAA litigation should be considered when developing disaster plans.

Other entities besides just covered entities can be implicated and at risk for PHI disclosure penalties. The "linked liability" concept may extend HIPAA-relevant liabilities and penalties to others. Covered entities forced into litigation for PHI privacy or security breaches may try to share blame or to pass it to others, such as the communications provider used to connect covered entity offices. Such other parties may also try to spread blame to yet others, such as negligent vendors of IT products with implementation errors, vulnerabilities, and less than adequate security. The ultimate assignment of liabilities may be legally tortious.

To minimize liability exposure, firms linked in even obtuse ways to covered entities are wise to be able to show (a) they exercised due diligence in establishing normal and customary privacy and security measures within their enterprises, services, or products; (b) their measures are in line with industry best practices and HIPAA requirements, and

(as required by HIPAA) anticipate and address new threats; and (c) they made correct decisions regarding their privacy and security measures.

Some argue that employers offering health insurance also have implicit liabilities linked to HIPAA. It is argued that employers have a duty to ensure healthcare plans and associated care providers offered to employees are compliant with HIPAA. Second, when insurance carriers send enrollment and benefits data to employers, employers may be construed as acting as business associates subject to provisions that business associates of covered entities must also be HIPAA compliant.

A further class of entity not considered *covered* under HIPAA is the property and casualty insurance carrier. One can argue that even though the insurance carrier does not dispense treatment and is not in the treatment chain, it handles the same information that HIPAA regulates. Many make a successful argument in favor of this special class in that the patient relinquishes any expectation of privacy when the incident (i.e., accident) becomes part of a suit or arbitration. In that case the information becomes public record. However, property and casualty insurance carriers may simply pay as underwritten with no suit or arbitration. This constitutes a sufficient gray area in covered versus noncovered entities.

HITECH updates to legacy HIPAA enforcement regulations appear in Section 71.5.6.4.

71.5.6 HITECH's HIPAA Update. The HITECH Act broadens and expands the scope of personal healthcare information (PHI) security and privacy beyond that stipulated by the original seminal HIPAA mandates. Changes to HIPAA regulations pertaining to healthcare information privacy, security, breach notification, and violation enforcement were recently issued in a so-called HIPAA Final Rule.[46] Of significance, the final rule now extends HIPAA's security and privacy requirements, HIPAA's compliance/penalty mandates, and liability risks to most business associates of HIPAA-covered entities.

71.5.6.1 HIPAA Security Enhancements. HITECH extends the original HIPAA Security Rule to apply full-force to most of the business associates of the covered entities to which the original HIPAA Security Rule applies. The business associates that are affected are those who are exposed to any PHI in the covered entity's environment.

Business associates must comply with all aspects of the HIPAA Security Rule. They must develop, implement, and maintain security policies and address the same security requirements as the covered entity. Healthcare entities must conduct due diligence in evaluating the information protection capabilities of prospective business partners before business relationships are established and formally consummated by formal business associate agreements.

The business associates impacted are those corporate or individual bodies that in the course of their business with a HIPAA-covered entity create, receive, maintain, transmit, or have access to protected health information entrusted to the HIPAA-covered entity. Impacted business associates include such diverse entities as, for example, vendors of PHI records, providers of claims processing, medical equipment vendors, medical transcriptionists, benefits managers, lawyers, cloud service providers, billing, accounting, actuarial, and legal and financial services, as well as other kinds of contractors and subcontractors that perform services on behalf of a healthcare-covered entity.

Current legal analysis[47] expects that impacted subcontractors of business associates will also include the service providers and technology companies that service the

business associates and that have access to PHI through their work for business associates as customers of their services or products. As such, HITECH stipulates that business associates will need to establish another tier of business associate agreements, namely business associate agreements with their subcontractors.

Further HHS guidance has been issued on how to secure protected health information.

71.5.6.2 HIPAA Privacy Enhancements. The HITECH enhancements to HIPAA privacy provide greater protection and control of PHI. HITECH extends all the use and disclosure provisions and limitations of HIPAA's original Privacy Rule to apply full-force to business associates of the covered entities to which the original HIPAA Privacy Rule applies.

Patient rights have been enhanced; patients have further control of the dissemination and use of their own PHI. For example, under specific patient payment conditions patients can prohibit providers from sending patient treatment information to health plan entities and insurance carrier entities. Patients can prohibit or authorize the sale of their PHI, the use or disclosure of their PHI for marketing and fundraising purposes, and/or the use of their PHI by researchers. Patients can authorize that certain of their minor children's PHI be provided to their children's school. Patients have the right to be notified of certain types of breaches of their PHI.

Providers' legacy privacy practice notices that were established in response to the original HIPAA Privacy Rule must be updated and posted to reflect the new HITECH HIPAA privacy mandates. All new patients must be given copies of such updated privacy notices.

71.5.6.3 HIPAA Breach Notification Enhancements. HITECH [48] and the FTC[49] updated and extended the breach notification requirements of the original HIPAA Breach Notification Rule. The original breach notification rule was based on the notion that breaches can pose a risk of harm to the individuals whose PHI was compromised. Such a risk of harm consideration was consistent with breach notification laws in several states. The new breach notification rule is based on the notion that there is risk that PHI is compromised instead of there being risk of harm to the individual.[50] The change to an *access* standard from a *harm* standard is intended to provide more objectivity in when to declare an incident to be a breach.

Furthermore, the applicability of breach rules now extends to the business associates of the covered entities to which the original HIPAA Breach Notification Rule applies. Business associates now have the same HIPAA liabilities and exposures as HIPAA-covered entities. As such, business associates must establish appropriate breach response policies and must respond to breach incidents in ways that are in full compliance with HIPAA regulations.

Legal analysis[51] of the extended breach notification requirements reveal that prudent covered entities and their business associates should move from a "risk of harm" standard to a "presumption of breach" standard. More specifically, parties impacted by HIPAA will move away from the original HIPAA position that a party that is the source of breached information is presumed innocent of harming patients when a breach occurs, until proven otherwise. They will move to the new HIPAA position that parties are presumed guilty of harming patients during a breach, unless the parties can prove their innocence. Impacted parties need to alter their incident response and breach notification strategies and processes, accordingly.

The new HITECH HIPAA requirements also modify what is classified as a breach and stipulate which breaches are reportable, under what conditions breach notifications must be issued, and to whom breaches must be reported.

The new requirements modify what risk assessments are required in order to initiate a breach incident response. Under the original HIPAA breach rules all unauthorized accesses to PHI were deemed breaches if such accesses posed significant risk of any financial or reputation consequences, or other type of harm, to the affected individual whose PHI was compromised. Under the new breach regulations, any access, disclosure, or use of PHI that is not allowed under the HIPAA Final Privacy Rule is a breach that must be reported unless the caretaker of the PHI in question can demonstrate that the probability of PHI exposure, via unauthorized PHI access or impermissible use by knowledgeable insiders, is too low to warrant a breach incident response.

The probability of PHI compromise demonstration must use a specific four factor risk assessment that must include factors like (a) whether the PHI in question was just viewed or whether a breach perpetrator actually acquired the PHI in question, and (b) the level of risk reduction brought about by efforts to mitigate the impact of the PHI exposure. Risk can be appropriately reduced, for example, by obtaining documented assurances from the breach perpetrator that the breached PHI has been destroyed, will not be further used, and/or will not be further disseminated.

Breaches that intentionally expose *unprotected,* personally identifiable PHI must be reported via breach notifications. Exposure of PHI that are *protected* by security such as encryption, but are nonetheless still compromised, do not require breach notifications.

Notifications of breaches affecting at least 500 individuals are to be conveyed to all individual victims as well as to HHS. Covered entities and their business associates have differing notification requirements that depend on which party discovers the breach. Notifications must be conveyed to at least each individual victim of unprotected PHI exposure, and possibly to other parties, within 60 calendar days after breach discovery. Only specified types of notification processes are permitted.

71.5.6.4 HIPAA Enforcement Enhancements. HIPAA enforcement actions are triggered by a written personal complaint that an HIPAA-covered entity or Business Associate violated some person's PHI privacy rights or otherwise committed some other violation of the HIPAA privacy or security rules. HITECH extended HIPAA civil and criminal penalties. Penalties for violations of noncompliance with any of the HIPAA Final Rules are now based on the level of negligence that occurred. The liabilities for Business Associates changed; the entire set of revised penalties, as well as HHS OCR audits and enforcement actions, is extended to apply full-force to Business Associates.

While business associates must now follow all HIPAA rules and can be held liable for an impermissible use or disclosure of PHI, the covered entities that receive services from their business associates are, nonetheless, still responsible for the actions of their business associates. As such, the chain of contracts—so-called Business Associate Agreements—from covered entities to business associates to business associates' subcontractors and so on must properly include provisions for appropriate assurances of HIPAA compliance.

Covered entities and their business associates must comply with all requirements stipulated in the HIPAA Final Rule by September 23, 2013. Any existing Business Associate Agreements between covered entities, business associates, or business associates' contractors or subcontracts that were in effect prior to January 26, 2013, must be modified by September 22, 2014, to reflect compliance with the new security, privacy,

and breach notification rule changes. The new enforcement rules apply after September 23, 2013.

According to HITECH, State Attorneys General are given the power to undertake HIPAA civil enforcement actions for HIPAA violations that impact their state's residents. The AGs can now seek damages on behalf of state residents.

HITECH also specifies that criminal and civil violations of HITECH and HIPAA mandates can be punished by a four-tier set of increasingly onerous penalties based on increasingly greater levels of culpability.

Criminal violations of HITECH HIPAA mandates are prosecuted by the U.S. Department of Justice and can result in imprisonment as well as financial penalty. In 2005, the DOJ weakened the original HIPAA criminal statue by claiming that this statue does not apply to individuals. Some have argued that such a weakening allows miscreants (e.g., hackers who steal healthcare information or covered entity staff who might sell PHI), to walk away free of consequences because they do not violate HIPAA.

HITECH specifies that three types of specific violations are prosecutable. The prosecutable violations include inappropriate use of a medical ID, theft of PHI, or disclosure of PHI. Contrary to the seminal HIPAA strategy of there being a single level of punishment, HITECH now specifies that persons knowingly engaged in prosecutable violation activities can be assessed any of several different penalties. For example, they can be fined up to $50,000 and be imprisoned for up to one year. If the intent of disclosure is to sell or use PHI for commercial advantage, imprisonment can be as high as 10 years and fines may be as high as $250,000. Such HITECH criminal penalties are over and above any criminal penalties resulting from violation of the seminal HIPAA Law.

Financial penalties for civil violations of HITECH depend on the nature and intent of the PHI disclosure. Contrary to the seminal HIPAA provision, HITECH specifies three levels of penalties. They range from a $100 to $50,000 fine per unknowing violation (with a $25,000 to $1,500,000 yearly cap) up to a $10,000 to $50,000 fine per willful neglect violation (with a $250,000 to $1,500,000 yearly cap).

The consequences of noncompliance to the HIPAA Rules are many, not just possible federal penalties and state fines. Business associates can suffer the loss of contracts. Covered entities and business associates can expect to cover the costs associated with creating and sending notifications to all individuals whose PHI was compromised. They may have to cover the costs of credit monitoring services for victims of information exposure. They may have to incur costs to repair hacked IT. They may have to incur costs in revising existing, or in establishing new, protection policies, procedures, awareness training, and backup capabilities. There will be numerous legal and security review costs associated with supporting criminal or civil investigations. There may be legal costs and possible monetary settlements of private legal actions by individuals impacted by their PHI loss or compromise—especially if compromised PHI was unknowingly altered and led to erroneous medical decision making. Loss of reputation may impact future earnings.

71.5.7 Demonstrating HIPAA Compliance. Establishing whether a healthcare entity is compliant with HIPAA rules can be a daunting task. Determining compliance with all the other laws and regulations applicable to a healthcare entity complicates the compliance picture even more. Just for the case of HIPAA alone, an organization must initially determine many factors. It must identify which of its care and business activities are under the jurisdiction of HIPAA. It must find, identify, and document the totality of PHI created, stored, processed, and displayed in the organization. It must identify PHI flows, both those internal within the organization as well as external flows,

to/from business associates and other third parties. Only then can an organization begin the process of assessing whether HIPAA rules are being followed in all situations.

As detailed elsewhere,[52] observations of the HIPAA compliance status of different healthcare entities in the healthcare sector show a mixed state of affairs. Large care providers (e.g., hospitals and medical centers) generally are complying in large numbers. The compliance story for other entities is mixed. Small care providers (single doctors and small offices of care providers) tend to be overwhelmed and minimally compliant. Insurance carriers are often not changing their legacy systems and procedures to adapt to HIPAA. Clearinghouses are trying to provide interfaces between providers and carriers, but these intermediaries are having troubles too.

Compliance-related initiatives associated with HIPAA-covered entities and business associates of covered entities are considered in Sections 71.5.7.1 and 71.5.7.2.[53]

71.5.7.1 HIPAA-Covered Entity Compliance. There are a number of organizations that support HIPAA compliance efforts. The organizations include the HHS Office of Civil Rights (Section 71.5.7.1.1), the healthcare accrediting commissions (Section 71.5.7.1.2), and the HITRUST industry group (Section 71.7.1.2).

71.5.7.1.1 HHS Office of Civil Rights (OCR). As required by HITECH, the HHS OCR established a compliance program for the HIPAA Privacy, Security, and Breach Notification Rules. The list of documents that may be requested by OCR during a HIPAA security onsite investigation and compliance review is extensive.[54] Compliance is assessed either by (a) OCR spot-checks of healthcare entities or (b) OCR targeted investigations triggered by reported incidents of PHI compromise.

71.5.7.1.2 Spot Checking. OCR must conduct spontaneous, spot-check, semi-random audits of healthcare-covered entities (and eventually business associates) to measure, against specified criteria, their compliance to the HIPAA Privacy Rule, Security Rule, and Breach Notification Rule. The spontaneous, spot-check audit is intended to result in a compliance improvement process whereby any deficiencies that are exposed can be fixed by the audited organization. Spot checks are repeated over time in order to assess changes in a healthcare entity's compliance status. Compliance status changes may arise from changes in healthcare business and care processes, personnel changes, HIPAA awareness changes, fielded technology changes, and so on.

However, it is possible that a spontaneous audit could unearth significant compliance shortcomings that could initiate a separate, targeted, OCR enforcement investigation (see below). The possible consequences of noncompliance discovered during a targeted investigation are many (see Section 71.5.6.4). As such, the threat of a spontaneous OCR audit may help[55] to improve an organization's compliance with HIPAA rules.

A pilot program to help OCR establish its HIPAA auditing role was recently completed. A wide range of large and small covered entities ranging from those using complex IT to those with paper-based workflows were included, including healthcare plans, clearinghouses, and individual and organizational providers. Each audit was completed within 30 business days. Each audit included onsite interviews and examinations of staff, equipment, facilities, and documents. Common shortcomings discovered[56] in this pilot included inadequate risk analyses, outdated HIPAA-related policies and procedures, and lack of contingency plans.

71.5.7.1.3 Targeted Investigations. OCR also conducts complaint-driven, targeted examinations (see Section 71.5.6.4) of specific covered entities purportedly

suffering PHI loss incidents, as well as business associates (see Section 71.5.7.2), that purportedly exposed PHI. These examinations are HIPAA enforcement investigations that establish the level of compliance in place at the time of an incident of PHI exposure. Investigations also assess whether a specific incident was truly a breach, what intent (e.g., unsuspecting versus malicious) was associated with the breach, and whether any appropriate sanctions and penalties are warranted.

71.5.7.1.4 Healthcare Accrediting Commissions. Electronic Healthcare Network Accreditation Commission (EHNAC) and the Utilization Review Accreditation Commission (URAC) are healthcare accrediting organizations establishing roles in being federally recognized, independent third parties for demonstrating or testing HIPAA compliance.

With the advent of possible onsite visits from the OCR and the increasing risk of exposure to HIPAA penalties and breach handling costs, healthcare-covered entities are looking to get their HIPAA privacy and security ducks in order before an OCR audit commences. An EHNAC or URAC privacy and security accreditation can position the accredited healthcare-covered entity with more evidence to complete the OCR audit successfully.

EHNAC is an outside resource that can be used by a healthcare-covered entity to support its privacy and security compliance both to HIPAA and HITECH.[57] EHNAC can be contracted by a healthcare-covered entity to evaluate its privacy and security policies, procedures, and implemented controls that may likely be audited by OCR for HIPAA compliance. EHNAC is also looking to position itself as a third-party security and privacy certifier of healthcare entities seeking compliance to the Direct Messaging aspect of HITECH Stage 2 Meaningful Use.

Similarly, URAC operates HIPAA privacy and HIPAA security accreditation programs.[58] These programs are used by HIPAA-covered entities and business associates to help demonstrate their HIPAA compliance. Like the comparable EHNAC process, the URAC HIPAA privacy and HIPAA security accreditation processes consider industry best practices and examine a healthcare organization's operational policies and procedures to verify effective compliance to HIPAA rules.

A HIPAA Privacy and Security Accreditation by URAC or EHNAC demonstrates a healthcare organization's good faith efforts to abide by HIPAA requirements. Receiving such an accreditation also provides a measure of HIPAA compliance assurance to a healthcare organization's prospective business associates. Patients can also get comfort that their PHI will be safe at an accredited healthcare organization.

Other private organizations, not necessarily federally recognized, also provide HIPAA compliance auditing services that evaluate whether and why a healthcare organization is HIPAA compliant and if not, why not.

71.5.7.2 HIPAA Business Associate Compliance. According to HITECH, technology service providers who want to make themselves available as candidate business associates to HIPAA-covered entities, as well as to other HIPAA business associates, can become certified that they are compliant with HIPAA. The HIPAA compliance certification ensures that candidate business associates comply with all of HIPAA's privacy, security, and breach notification requirements so that a contracting HIPAA covered entity, or any contracting HIPAA business associate, is confident that when their PHI is exposed to a compliance-certified business associate, the PHI will be protected to the same level specified in HIPAA and that the compliance-certified business associate's protection solutions are HIPAA-compliant. The compliance-certification also ensures

that candidate business associates conduct recurring evaluations that demonstrate that over time their privacy, security, and breach notification policies and procedures and associated awareness training meet all of HIPAA's privacy, security, and breach notification requirements.

71.5.8 Related Federal Guidelines. Various departments and agencies within the U.S. Government develop guidelines and other materials pertinent to the security and privacy within the healthcare sector. General guidelines are also available that apply across any industry or sector.

For example, HHS makes available several online healthcare security and privacy regulation summaries and guidelines, privacy and security tools, awareness and training materials, brochures, and videos.[59] Topics include, for example, providing guidance for planning implementation of solutions meeting HIPAA and Meaningful Use Privacy and Security requirements[60]; providing education about HIPAA-stipulated risk analysis[61]; providing guidance on using the HIPAA Privacy Rule when electronically exchanging PHI[62]; giving HIPAA Security Rule implementation support[63]; providing guidance on the remote use of PHI[64]; and providing guidance on protecting PHI.[65,66]

The National Institute of Standards and Technology (NIST) developed a number of white papers that provide insight into the HIPAA security rules.[67] NIST also published guidance to assist in understanding, implementing, and evaluating operational implementations of HIPAA security requirements.[68,69]

HHS also provides guidelines that specify basic steps for securing mobile devices that need to protect healthcare information.[70,71] The National Institute of Standards and Technology (NIST) has additional recommendations[72,73] for mobile devices. The NIST material is applicable across all fields of endeavor.

The Federal Trade Commission has developed guidance pertaining to several aspects of protecting healthcare information, including medical identity theft, risks of electronic file sharing, and electronic digital copiers.[74]

71.5.9 Comparing Federal and State Healthcare-Relevant Security and Privacy Laws. HIPAA provides a national floor for requirements for protecting PHI; state and other federal laws may provide even stronger protections. Examination of the mosaic of state laws introduced in Section 71.6 reveals that the existing state laws pertaining to PII are to some degree overlapping and sometimes inconsistent in terms of the enacted PII protections. In some cases, HIPAA may enhance the protections afforded by other laws. In other cases, other laws may deal with privacy and security more stringently than HIPAA.

So-called HIPAA preemption analyses are available in many states[75] to understand the relationship between state-specific privacy law and the HIPAA Privacy Rule. In general, the stronger information protection law should prevail. However, conflicts between state law and HIPAA are not always straightforward to resolve, especially when there are differences in limitations or restrictions that are, or are not, precisely stipulated in state and federal laws. Sometimes judicial ruling is needed.[76]

As the federal desire to implement a nationwide health information infrastructure based on sharing of person-specific EHRs progresses, differences among the state-specific laws need to be reconciled. The Harmonizing State Privacy Law Collaborative was formed as part of the Health Information Security & Privacy Collaboration (HISPC) to assist in such matters. The HISPC also considers federal–state healthcare coordination issues related to (a) harmonizing state-specific healthcare information release consent, (b) developing interorganization agreements pertaining to electronic

healthcare information exchange, and (c) harmonizing certain aspects of security to facilitate secure PHI exchange between differing Health Information Exchange (HIE) models used in different states.[77]

The precedence between privacy and security requirements of state laws and the federal HIPAA law can have problematic consequences. Interoperability issues can arise between parties that have stronger protection methods and those that have the common, minimal level of protection stipulated by HIPAA. As such, there is the potential for problems with the concept of nationally networked EHRs, as described later in this chapter. Development and exchange of electronic healthcare records can be hampered by inconsistencies and variations in privacy and security policies and regulations between jurisdictions. The HISPC detailed variations in state laws and policies that pertain to the privacy and security of health information that is electronically exchanged.[78] Furthermore, there may be legal questions regarding how patient data can be shared across state lines. Some argue that HIPAA needs to be modified to provide more comprehensive national information protection standards that form a consistent, mandatory uniform ceiling above the growing hodgepodge of similar but potentially inconsistent state laws.

Exhibit 71.2 provides a comparison of the differing breach notification requirements in a representative example of three different laws: two federal laws (GLBA and HIPAA/HITECH) and one state law (Massachusetts). Covered entities have the burden to prove which law or laws pertain to each breached, individually identified data item. It is possible that a breached entity will be held accountable by several laws. Each law runs simultaneously. As such, a covered entity must respond to each breach law and in accordance with each law's timelines and other breach notification requirements.

Breach complaints, whistle blowers, or voluntary admission drive the enforcement of state breach laws and likewise the enforcement of federal breach regulations. Unlike the federal government, states typically do not have the resources to conduct active state-run policing through audits. Often, federal penalty assessments take precedent because states stand back as federal penalty actions occur. Sometimes, states still make penalty assessments that are additive to federal penalties.

Exhibit 71.3 provides an example comparison of a specific set of the privacy requirements appearing in many different laws. In particular, the exhibit identifies which specific date elements of PII have privacy requirements in the various laws. In particular, Table C1 identifies the eighteen primary HIPAA PHI data elements and compares them to the personal data elements identified in PCI DSS 2.0, GLBA, state of California breach notification, and the commonwealth of Massachusetts' 201 C.M.R. 17.00. Initially evident is that data included in HIPAA as PHI is far more inclusive. Regardless of this scope change, those practitioners charged with PHI care may note that there are many instances where data elements overlap.

In lines 1 and 2, the element "name" applies to all drivers listed. California, Massachusetts, GLBA, and PCI all specify a level of detail, whereas HIPAA does not. Using HIPAA's broad scope as a cue, it can be assumed that "names" includes either first or last name rather than the combination denoted in all other statutory and regulatory guidance listed. Of course, a name alone and outside the context of health information (or PII, PFI, PCI) is just a name. However, within the context of a healthcare record, even a first name requires protection.

For the practitioner charged with protecting healthcare information, the nuances of HIPPA data elements are important. Though, many of the controls already applied in other areas simply apply to a greater scope of data elements thus making the problem less overwhelming.

EXHIBIT 71.2 Comparison of Requirements in Various Breach Notification Laws

Requirement	GLBA	Massachusetts 201 C.M.R. 17.00	HIPAA/HITECH
Notification Exception	When no misuse of data has occurred or when reasonably assured that none will occur (e.g., encryption with proper key control).	If encrypted electronic data and the confidential process or key that is capable of compromising the security, confidentiality, or integrity of personal information not in question	Unintentional, but good faith access by employee or individual acting under covered entity that does not lead to further disclosure, inadvertent disclosure to otherwise authorized individuals does not constitute a breach and disclosure/notification not required. Additionally, breach must also actually compromise the security or privacy. Ostensibly encrypted information with adequate key management would not compromise security or privacy.
	Good encryption with adequate key control usually provides an exception to notification for electronic records. However, paper records have no such exception.		
Notification Timeline	Notify when an institution becomes aware of unauthorized access and misuse of data has occurred or reasonably possible	As soon as practicable and without unreasonable delay, when the Entity knows or has reason to know of a breach of security, or when the Entity knows or has reason to know that the PI of such resident was acquired or used by an unauthorized person or used for an unauthorized purpose.	Discovered is the first day the breach is known or when the entity should have known. From this point there should be "no unreasonable delay and in no case later than 60 days." Sec 13402 (c).
	There is no minimum threshold that changes the method or timing of notification.	HIPAA states that if breach affects 500 or more individuals, without unreasonable delay and in no case later than 60 days from discovery. Additionally, if breach affects less than 500 individuals, within 60 days of the end of the calendar year in which it occurred. HITECH does not make the size distinction. Rather HITECH requires no unreasonable delay and in no case later than 60 calendar days after discovery.	
	All provide for reasonable "law enforcement" delays to protect criminal investigations.		

Notification Recipients	To customer in a manner in which the customer can be reasonably expected to receive it.	MA Resident Affected	Written notification to individual (next of kin if individual is deceased).
		MA Attorney General	If contact information is insufficient or out of data for 10 or more individuals, prominent Web notification, local print or broadcast media shall be used.
		Director of consumer affairs and business regulation (who may notify others as necessary)	Notice to Secretary of HHS (Web form), Posting to HHS Public Website.
Notification Information	Description of incident including the kind of customer information compromised, what the institution has done to protect from further unauthorized access, contact information (e.g., telephone number) for assistance, a reminder to remain vigilant for 12 to 24 months, instructions to review and report any suspicious account activity, description of fraud alerts and how to place them in consumer credit reports, reminder to obtain credit reports and have fraudulent activity removed, and explanation of how to retrieve a free credit report.	Resident: residents affected, and what actions the covered entity is taking to address the incident	Brief description of breach, date (if known), types of unsecured information (e.g., name, SSN, account number, etc.), steps individual should take to protect themselves, what the covered entity is doing to investigate the breach, mitigate losses, and prevent further breaches, and contact procedures.
	Attorney General & Director of Consumer Affairs: Above and Nature of the breach or unauthorized access or use		
	A common theme is that notification must include what the entity is doing currently to prevent further breaches.		
Breach Applies to Any Form Including Paper	Yes	Yes	Yes
	Breach notification most commonly applies to electronic records. However, most breach notification does not discern between paper records and electronic records.		

EXHIBIT 71.3 Comparison of Requirements in Various Privacy Laws

Data Element	HIPAA PHI	CA PII	MA PII	GLBA PII/PFI	PCI DSS PFI/PCI
1st Initial or Name & Last Name	Yes	Yes	Yes	Yes	Yes
Names	Yes			Must include first initial or name AND the last name whereas HIPAA has no such distinction.	
Account Number, Credit Card Number, Debit Card Number, with any required security code, access code, or password that would permit access	Only if it is associated with health care information	Yes	Yes	Yes	Yes
Account numbers	Yes			Only if the number is associated with personally identifiable information or personal financial information.	
All elements of dates (except year) for dates directly related to an individual, including birth date, admission date, discharge date, date of death; and all ages over 89 and all elements of dates (including year) indicative of such age, except that such ages and elements may be aggregated into a single category of age 90 or older	Yes	No			
All geographical subdivisions smaller than a state, including street address, city, county, precinct, ZIP code, and their equivalent geocodes, except for the initial three digits of a ZIP code if, according to the current publicly available data from the Bureau of the Census: (1) the geographic unit formed by combining all ZIP codes with the same three initial digits contains more than 20,000 people; and (2) the initial three digits of a ZIP code for all such geographic units containing 20,000 or fewer people is changed to 000.	Yes		No	This is partially applicable to GLBA and recent court cases indicated that postal code is private information in the context of a credit card transaction if not required as part of the transaction.[1]	

EXHIBIT 71.3 (*Continued*)

Data Element	HIPAA PHI	CA PII	MA PII	GLBA PII/PFI	PCI DSS PFI/PCI
Any other unique identifying number, characteristic, or code (note this does not mean the unique code assigned by the investigator to code the data)	Yes	No			
Biometric identifiers, including finger and voice prints	Yes	No			
Certificate/license numbers	Yes	No			
Device identifiers and serial numbers	Yes	No			
Electronic mail addresses	Yes	No			
Fax numbers	Yes	No			
Full-face photographic images and any comparable images	Yes	No			
Health plan beneficiary numbers	Yes	No			
Internet Protocol (IP) address numbers	Yes	No			
Medical record numbers	Yes	No			
Phone numbers	Yes	No			
Social Security Number	Yes	Yes	Yes	Yes	
Medical and/or Health Insurance Information	Yes	Yes		Yes	
Vehicle identifiers and serial numbers, including license plate numbers	Yes	No			
Web Universal Resource Locators (URLs)	Yes	No			
Information in Cookies	Not specifically listed; however, any information here in the context of healthcare information could be subject to regulation	No		Yes	No
Information on Consumer Reports	No	Yes	No		
State Driver's License Number	Yes	Yes			
State ID Number	Yes	Yes			

[1]Tyler v. Michaels Stores, Inc., No. 11-10920, 2012 WL 32208 (D. Mass. Jan. 6, 2012).

Operationally, the impact of a healthcare entity's need to comply with multiple laws and regulations arising from many jurisdictions can result in many complicating factors that a healthcare entity's information protection professionals must consider when managing PHI in a real operational environment.

The scope of HIPAA data elements extends beyond the boundaries of PII, PFI, and PCI. This only means that controls employed to protect other data elements necessarily apply. Erring on the side of caution by employing broader controls is prudent. Additionally, applying discrete controls in the absence of a data classification standard and risk-management program introduces complexities to an already complex arena. However, application of NIST 800-53 (see Section 71.6.2.2), or other risk frameworks (e.g., ISO 27001) generally provides adequate organization to address healthcare information security. There are pros and cons to each method and, like much guidance in information assurance; it must match the organization's culture and ability to execute given business and resource constraints.

There is immense value in information classification as an abstraction exercise. That is, classification, implemented as a means to determine the value of data and to identify the most appropriate way to store, to manipulate or to access data, brings immediate management value. It establishes classes and categories of like data for similar treatment. Classes establish a label whereby the healthcare entity can set policy, determine action, and protect assets based on known drivers. In effect, information classification can determine the minimum level of protection required based on the nature of the classification of the data. Furthermore, classification provides knowledge and guidance to determine the most effective course of action.

Where classification reveals that current systems, applications, processes, and procedures present risk, then risk remediation action is possible. Action will range from acceptance of risk, to implementation and fielding of additional controls, to redesign of IT systems or security protections. In the context of healthcare data, where controlled data have a wider scope than other statutory and regulatory edicts, classification serves to ensure that both security practitioners and nonsecurity employees understand data class, the driving regulation (e.g., HIPAA, PCI, etc.), and the required protection treatment (e.g., encryption at rest, in flight, or both, etc.).

As a practical matter, managing to compliance with HIPAA or any other statutory or regulatory driver diverts focus to the act of compliance rather than the act of managing risk associated with the data. HIPAA, as a driver, is one of many within the context of risk management within a healthcare entity. This is not to lessen its importance, but healthcare entity protection professionals must determine the most stringent of driver requirements and manage to that baseline. Each healthcare entity must write this most stringent baseline into its policies, standards, procedures, and enterprise awareness training. The minimum requirement becomes that of the entire enterprise and sufficiently covers all drivers, in all locations, and at all times.

71.6 OTHER PUBLIC SECTOR HEALTHCARE PRIVACY AND SECURITY INITIATIVES.

There are several government-instigated initiatives relevant to healthcare information privacy and security. HITECH efforts at pushing the United States toward so-called "meaningful use" of a new healthcare model that relies on secure, nationwide, electronic sharing of protected healthcare information are considered in Section 71.6.1. The several information protection regulations and guidelines that apply to DoD healthcare are considered in Section 71.6.2. Yet other government initiatives related to healthcare information protection are identified in Section 71.6.3.

71.6.1 The HITECH Act: Moving to Electronic Health Information. In part following on from a 2004 Presidential Executive Order,[79] the 2009 Health Information Technology for Economic and Clinical Health (HITECH) Act[80] funded and tasked the U.S. HHS to expand the adoption of health information technology and to create a national network of secure and private electronic health records (EHRs). In conjunction with its HITECH efforts, HHS also updated and expanded the HIPAA Rule of 1996 (see Section 71.5.6).

HITECH established an Office for a National Coordinator (ONC) for Health Information Technology (HIT). ONC established and guides the development of national HIT policy, standards, and certification testing associated with a national HIT infrastructure supporting the electronic use and secure exchange of privacy-protected EHRs. While ONC is currently incentivizing physicians and hospitals to adopt and to use certified EHRs, in 2015, healthcare entities and providers will receive Medicare financial penalties if EHRs are not used.

The ONC Meaningful Use (MU) initiatives are focused on the staged development, incentivized adoption, and meaningful use of protected EHRs. The MU initiatives are summarized in Section 71.6.1.1. The privacy and security aspects of the various stages of MU are described in Section 71.6.1.2.

71.6.1.1 HITECH's Meaningful Use Initiatives. Traditionally, the primary information that was exchanged among different parties in the healthcare sector was the exchange of batched billing information between providers and carriers, as well as the sending of prescriptions from providers to pharmacies. The primary way of communicating healthcare-related information among healthcare parties was fax exchanges of unstructured data.

HITECH's MU initiative[81] puts the U.S. healthcare system on a path leading to a nationwide network of interoperable, patient-specific protected EHRs. The EHRs will contain structured and unstructured medical information that can be electronically and securely processed, exchanged, and used by any care provider in the country. EHRs will contain all pertinent patient PHI. Subject to privacy constraints, trusted EHRs can be used to coordinate and to direct any patient's care among any set of care providers anywhere in the country.

MU's EHR efforts are looking to shift the healthcare delivery sector toward a new paradigm wherein a patient and all her/his applicable care providers get complete and secure visibility into all *appropriate, need-to-know,* private medical information of the patient. "Appropriate" information is the patient information applicable strictly and only to the specific patient conditions being treated by, and the services being provided by, each of a specific patient's various specific care providers. Furthermore, it will be possible to know who has accessed any particular patient data element.

The MU strategy is to move though various and expanding stages of capabilities that in aggregate build up to the ultimate goal of nationally interoperable and appropriately protected EHRs. Currently, three stages are specified.

71.6.1.1.1 Stage 1 MU. The goals of Stage 1 are to get providers and hospitals (a) to collect specified, relevant patient data electronically so as to be able to create EHRs, and (b) to use EHRs to meet a specified set of healthcare service requirements. Specific metrics were defined to evaluate whether fielded EHR technology met Stage 1 MU and, therefore, that healthcare entities could be awarded government incentive payments for fielding certified EHR technology.

71.6.1.1.2 Stage 2 MU. Stage 2 MU expands requirements on EHRs and creates the concept of Health Information Exchanges (HIEs). Again, metrics were established to define what it takes for a healthcare entity to claim that their fielded technology met Stage 2 MU.

The goal of MU Stage 2 is to start developing electronic capabilities to share, in real time and with sufficient security, protected clinical and patient care information as captured in an EHR. Such information sharing can increase the continuity of a patient's care among all providers seen by a patient. Stage 2 progresses through several steps toward demonstrated, bidirectional, more capable, interoperable, and secure exchange of ONC-specified standard patient records among different, disparate EHRs in disparate healthcare entity environments. Heretofore, interoperability among disparate EHRs was hampered because of environment-specific use of different vendor EHR products that create, store, display, communicate, or message EHR content data in different, sometimes proprietary ways.

Several federal, state, and industry groups[82] are working toward the Stage 2 goal of securely moving care data among disparate EHR systems. Interoperation can be via secure direct EHR to EHR interoperation, or it can be via an intervening Health Information Exchange. HIEs are trusted intermediary systems through which the data from the EHR system of one healthcare entity can move to the EHR system of another healthcare entity. States or regions within a state create and maintain state-specific HIEs. The aggregate of such HIEs are to form a Nationwide Health Information Network (NHIN)—aka the eHealth Exchange.

A certain portion of initially exchanged EHR data must consist of at least specific summary care records. Interoperation must be capable among and between several different groups. Such groups include (a) providers—both large (e.g., hospitals) and small (e.g., individual practitioners), (b) referred or referring providers, and (c) patients.

The strategy for ensuring interoperability relies on use of data transport, data structure, and data content standards[83] that are uniformly implemented by interacting parties. Interoperability is demonstrated and assured by formal EHR certification that arises from standard testing of EHR interoperation. Such testing is done in conjunction with a recognized, independent third party.

The focus on interoperation with patients is critical and essential, since a goal is to create the healthcare sector equivalent of a customer relationship management capability. But it is realized that electronic interoperation with patients and allowing patients to access their longitudinal health records or to send secure messages to care providers may be problematic. Certain patients may have less familiarity with Internet and computing technologies. As such, initial patient access to provider-created data may occur using air-gap technologies such as data exchange via the physical passing of a USB drive or a paper printout.

Of significance to this chapter, and as described below, all interoperations must be secure and private.

71.6.1.1.3 Stage 3 MU. The purpose of Stage 3 MU is to move the Medicare and Medicaid community to a new model of care made possible by meeting the Stage 2 goal of sharing patient EHRs. The new model is to be a collaborative model that focuses on care outcomes and shares care responsibility and care accountability among contributing providers, patients, and families. Where necessary the new model capitalizes on HIEs. The specific goals of Stage 3 MU, and the metrics to evaluate

achieving Stage 3 MU, were not yet established at the time of the writing of this chapter. Indeed, the group that advises HHS on health IT sought community inputs to shape the Stage 3, and later-Stage, MU objectives during 2013.[84]

71.6.1.2 HITECH's Meaningful Use Security and Privacy Initiatives. In the course of pursuing Meaningful Use in the exchange of PHI via EHR technology, new healthcare information protection issues arose that were not anticipated in the HIPAA Privacy and Security Rules.

For example, HIPAA does not accommodate the situation where (a) a patient authorizes the sharing of EHR, (b) only a patient-authorized, specified portion of the PHI from a specific patient's EHR needs to be shared among multiple parties when multiple parties are involved with the care of a specific patient, and (c) mutual trust in the confidentiality of, and any assigned redisclosure rules associated with, shared PHI by all parties needs to be assured. Accordingly, the ONC initiated the Data Segmentation for Privacy Initiative (DS4P) to support the privacy and security aspects—such as exchange, adjudication, and enforcement of privacy consents—associated with sharing EHR PHI across the domains of different healthcare entities.

The DS4P used collaborative members of ONC's Standards and Interoperability Framework[85] to examine the use of data segmentation[86] to accomplish privacy consent management. Data segmentation ensures that any elements of PHI that are deemed unsuitable to share are indeed not shared. Segmentation is supported by privacy metadata tags that convey the disclosure policies associated with each tagged date element within PHI. These tags specify (a) the confidentiality that needs to be maintained on the PHI data element being shared, (b) all obligations (e.g., nonredisclosure) of the party receiving the shared PHI data element, and (c) the purposes (e.g., emergency treatment only) for which the shared PHI data element are allowed to be used.[87]

To encompass the new healthcare information protection requirements as well as the requirements within the HIPAA Privacy and Security Rules, the ONC developed a framework[88] that establishes the principles associated with a consistent national approach to privacy and security for the storage and exchange of EHRs. More recently, ONC developed guidance pertaining to MU privacy and security.[89] The guide also enumerates a plan for meeting MU privacy and security requirements.

In relation to the HITECH MU initiative and the three defined MU stages, the ONC has developed specific MU security and privacy requirements that pertain to EHRs.[90] It is specified that an EHR must provide adequate security and privacy protections for PHI by meeting nine security requirements: (1) access control, (2) emergency access, (3) automatic log-off, (4) audit log, (5) integrity, (6) authentication, (7) general encryption, (8) encryption when exchanging electronic health information, and (9) accounting of disclosures. The ways by which such requirements are to be met are not mandated.

ONC uses specified standard, implementation specification, and certification criteria to establish the technical capabilities of EHRs and to certify EHR products.[91] The scope of certification criteria includes privacy and security. The ONC Health Information Technology Certification Program maintains a list of EHR technology that has been tested and certified by an authorized body for achievement of Meaningful Use.[92]

As outlined below, different security and privacy requirements are required to achieve different stages of MU.

71.6.1.2.1 MU Stage 1 Security and Privacy Directives. In Stage 1, patient data must be protected to assure its privacy. As per the HIPAA Privacy Rule, patients must

be able to obtain an electronic copy of their healthcare information. As per the HIPAA Security Rule, electronic PHI handled by EHR technology must implement policies and procedures for dealing with security violations. Healthcare entities looking to comply with Stage 1 must attest[93] that have met the specified HIPAA privacy and security requirements required for Stage 1.

Healthcare entities looking to comply with Stage 1 must engage in a continuing risk-management process. Periodically, one must conduct or review an EHR security risk-assessment, correct any security deficiencies that are found, and provide security updates.[94] While not mandated, possible ways of meeting such requirements are provided by the Centers for Medicare & Medicaid Services (CMS).[95]

In the field, however, astute healthcare parties are treating the risk analysis directive in various ways. Some see the risk analysis as nothing more than a be-all, end-all activity to meet the MU Stage 1 information protection directive. More astute parties use the risk analysis to identify, to understand, and to prioritize their healthcare information vulnerabilities and to initiate efforts to remedy discovered vulnerabilities.

71.6.1.2.2 MU Stage 2 Security and Privacy Directives. In Stage 2, complying providers and hospitals are required to attest that healthcare information at rest is encrypted.

In Stage 2, certified EHR products are required to include a secure, interoperable, email capability called the *Direct Secure Messaging protocol*. The Direct Project offers this email capability for point-to-point healthcare information exchange.[96] ONC's State Health Information Exchange Program runs the Direct Project.[97] Direct email has been demonstrated to exchange healthcare information among providers and hospitals in several states and across state borders.[98]

One possible Direct strategy for securing the privacy of information being messaged is use of a metadata wrapper around the message. The wrapper can stipulate message routing and the security needed to maintain privacy. Direct messages are transmitted via health information service providers (HISPs). HISPs are Internet Service Providers for the healthcare community.[99] They support the Direct email protocol and encrypt Direct messages. HISPs are registered with a registration authority and have digital certificates from a certificate authority. Authentication of Direct messages is done by the HISP. Vendors' EHR products must be certified in conjunction with one or more of such HISPs.

Another possible and more flexible Direct strategy relies on use of Web-focused SOAP (Simple Object Access Protocol) messaging services with digital signature, encryption, and authentication security services. SOAP-oriented security services are independent of HISPs and therefore do not require certification with specified HISPs or specified HIEs.

71.6.1.2.3 MU Stage 3 Security and Privacy Directions. The Stage 3 healthcare information protection goals were not yet stipulated at the time of the writing of this chapter. It appeared there was interest in considering issues of maintaining and managing healthcare information privacy and security, as well as consent information, as healthcare and consent information flows through HIEs from one domain of information protection capabilities to some other domain with potentially differing information protection capabilities.

A security goal that is likely to be incorporated into Stage 3 MU is that EHRs and providers should be able to conduct multifactor authentication, meeting NIST 800-63-1 Assurance Level 3,[100] for providers looking to access remote PHI outside the domain

of the provider's healthcare entity. Also under consideration for inclusion in Stage 3 directives are requiring attestation that (a) certain aspects of security awareness, and training, are performed, and (b) logs of authorized and unauthorized access to PHI are maintained.

71.6.2 Healthcare Information Protection in the U.S. DoD Healthcare Environment. Certain government agencies provide healthcare, or otherwise handle PHI. In particular, the U.S. DoD operates several hospitals and medical facilities that create, store, process, manage, and transmit personal healthcare information. The DoD also exchanges PHI among different facilities. It is also looking to transition to unified medical records that are shared across DoD as well as its Veterans Affairs programs and hospitals. As such, all HIPAA, HITECH, and MU rules, standards, and directives pertaining to security and privacy of healthcare information, to reporting of healthcare information breaches, and to the pursuit of breach enforcement actions must be upheld in the DoD healthcare sector.

However, all government agencies—whether civil, intelligence, or defense—must also comply with the Federal Information Security Management Act of 2002 (FISMA). Accordingly, any government agencies that deal with the creation, processing, storage, and transmission of healthcare information must comply with multiple regulatory laws and multiple regulatory authorities for the purpose of protecting healthcare patient information and healthcare business processes. Such agencies must comply with both HIPAA, as administered by HHS, OCR, and DOJ, and the Federal Information Security Management Act of 2002 (FISMA), as administered by object management group (OMG).[101] And, in certain instances such agencies can also be compliant to yet more regulations and standards such as, for example, the Federal Risk and Authorization Management Program (FedRAMP) and Common Criteria.

71.6.2.1 FISMA. FISMA requires all federal agencies to adopt and to demonstrate use of NIST-specified IT security processes and procedures that include categorizing risks associated with IT systems, developing and maintaining minimum controls to protect information systems, verifying and monitoring the effectiveness of their security controls via a specified certification and accreditation framework, and taking corrective actions when necessary.

71.6.2.2 NIST 800-53. FISMA requires that all Federal agencies, including Federal healthcare entities, as well as their contractors and service providers, must adopt and demonstrate use of NIST Special Publication (SP) 800-53, as well as other specified NIST Federal Information Processing Standards and Special Publications. NIST SP 800-53[102] establishes a unified information security framework that applies across the entire Federal Government. It specifies that parties that need information protection must use information assurance best practices. Such practices provide and demonstrate the due care and diligence necessary to fulfill their fiduciary responsibilities associated with being custodians of PII, in general, or PHI, more specifically.

The best practices[103,104,105,106] stipulated in NIST SP 800-53 include the following information protection processes and procedures:

1. Use of a specified risk-management framework and process to achieve appropriate information protection sufficient for the situation wherein information protection is needed;

2. Use of specified, customizable, policy-neutral, security[107] and privacy controls from extensive, periodically evolving catalogs of augmentable security and privacy controls:
   (a) to provide information-specific, technical, and operational information protection safeguards and countermeasures that meet the identified security and privacy requirements and the risks, threats, exploits, and vulnerabilities associated with the environment in which the security and privacy requirements exist, and
   (b) to protect organizational-level programs for managing information protection activities;
3. Use of specified, consistent, repeatable guidelines for selecting the best information protection controls to meet specific information protection requirements;
4. Assurance requirements for establishing confidence in the correct selection, implementation, and operation of the selected information protection controls;
5. Use of a catalog of customizable, extendable assurance assessment actions to stipulate actions that provide evidence that the risk-management processes and information protection controls used are correctly implemented, operate as intended, produce the desired outcomes, are otherwise effective, and may be worthy of receiving security accreditation; and
6. Procedures for periodically demonstrating at a point in time—and eventually continuously demonstrating (e.g., via a potential DHS Continuous Monitoring as a Service capability)—the confidence and effectiveness of the selected information protection controls meeting the identified information protection requirements.

The FISMA/NIST 800-53 approach has its detractors who (a) identify possible problems and costs of using single-point-in-time compliance-oriented strategies to provide information protection and (b) propose use of a smaller set of security controls. Having seen disconnects between FISMA compliance reports to the OMB and an agency's actual security posture, such individuals believe that compliance is a deception. They believe that information protection efforts must be solely focused on providing good security controls rather than on demonstrating point-in-time compliance to regulations.

Regardless of differing views, information assurance does improve when rigorous information protection processes and procedures are followed. It is reasonable to expect that a combined approach that depends on building security features correctly from the get-go (the NIST 800-53 strategy) and then continuously monitoring fielded products and systems will be an approach that most everyone can support.

71.6.2.3 FedRAMP. It is reasonable to expect that in time government agencies that provide healthcare, or otherwise handle PHI, will adopt use of contracted, cloud-based technologies and services. Such government agencies can choose to use the Federal Risk and Authorization Management Program[108] (FedRAMP) to address the security needs of PHI that is processed, stored, or exchanged in cloud-based environments.

FedRAMP extends FISMA and NIST SP 800-53 applicability to cloud computing situations. In particular, FedRAMP provides a government-wide, unified risk management program that (a) establishes unified, government-wide security requirements for cloud technologies and services, (b) authorizes the security capabilities provided by specific cloud computing products and services for government agency use, and (c)

uses accredited third parties to assess, via continuous security monitoring, the security protections provided by such products and services.

The IBX Vault (www.ibxvault.com) is an example of an external, third-party facility that can be used to provide contracted managed hosting, data center, and cloud services that are compliant to several government mandates including FISMA, the DOD Information Assurance Certification and Accreditation Process (DIACAP), FedRAMP, HIPAA, HITECH, PCI, and Sarbanes-Oxley (SOX).

71.6.2.4 Common Criteria. DoD's National Security Agency is party to a multicountry agreement to develop Common Criteria–specified profiles of security requirements and accompanying product implementation assurance requirements for specified types of commercial, off-the-shelf, IT products. A list is available of third-party–evaluated IT products that claim conformance to specific information protection requirements profiles and that have been independently validated[109] to comply with the claimed profiles. Such products were evaluated by any of several specifically certified evaluation laboratories that use a Common Criteria–specified standard international security testing and evaluation scheme. U.S. Government agencies that provide healthcare or otherwise handle PHI may choose to deploy such evaluated and validated products to meet those of their healthcare information security requirements that are consistent with the validated information protection solutions implemented by such products.

Historically, in the late 1990s NIST was looking to assist the healthcare community (both public and private), as well as the companion healthcare accreditation and certification community, by supporting a cross-healthcare-industry effort to develop profiles of HIPAA-based security and implementation assurance requirements of value to the healthcare community. Such requirement profiles were based on the international Common Criteria standards. The proposed strategy was for healthcare entities and vendors to use certified Common Criteria testing labs to verify the correctness and completeness of their HIPAA-based healthcare information protection solutions to specified HIPAA-based healthcare security and assurance requirements profiles. Accreditation and certification of healthcare entity use of such solutions would be straightforward as healthcare accrediting and certification agencies, such as EHNAC and URAC, would recognize and approve the use of such verified solutions.

This NIST-supported effort was not successful. Perhaps it was too early in the life cycle of healthcare information protection regulations. Apparently, healthcare corporate managers were not yet prepared to commit or, in their eyes, divert sufficient healthcare entity resources into healthcare information protection. They may have been reluctant due to the newness of government regulations and the unknown seriousness with which the government would police and enforce such regulations.

Perhaps now that the seriousness of widespread government regulation enforcement and the associated enormity of imposed penalties and fines are realized, the healthcare sector will consider revisiting such a standards-based approach to verifying the correctness and completeness of their information protection solutions.

71.6.3 Evolving and Emerging Government Healthcare Information Protection Initiatives. Many government bodies are pursuing, or seem poised to initiate, other security and privacy initiatives of relevance to the healthcare sector. Examples that existed at the time of the writing of this chapter appear below. It is not clear which—if any—of these examples will become fruitful initiatives.

The HHS ONC continues to pursue security and privacy initiatives related to HIPAA, HITECH, and Meaningful Use.

The Information Security and Privacy Advisory Board (ISPAB)[110] has lobbied for some Federal Government body such as the FDA to initiate a program to evaluate the security provisions of medical devices before they can be brought to market. Compromise of patient information and medical device control messages will likely become a big problem (see Section 71.7.2). The Federal Government may provide assistance in reducing security and privacy issues in the medical device field by establishing a program to provide some sort of "Seal of Approval" for successfully assessed medical device products.[111] Such a seal could alert buyers that specific products are approved for the security provisions they provide. How such a seal might differ from HITRUST-certified assessments or any Common Criteria certifications has not yet been addressed

The ISPAB is also pushing for the Feds to create a United States Computer Emergency Readiness Team for reporting and tracking security vulnerabilities within medical devices as well as security incidents associated with medical devices. Whether or not the government will undertake any ISPAB-requested initiatives is not known at this time.

In early 2013, the U.S. President issued an Executive Order[112] and a companion Presidential Directive.[113] This order tasks the NIST to develop a cybersecurity framework for protecting critical infrastructure in the United States. It is expected that this Executive Order and the ensuing activities will lead to new cybersecurity legislation. The United States considers the healthcare sector to be part of the U.S. critical infrastructure. Whether any dramatically new U.S. cybersecurity framework will emerge in the near term or whether any new legislation arising as a result of the NIST efforts will make changes to existing healthcare-related security and privacy laws and regulations is unknown. Readers are encouraged to stay abreast of any such new legislation.

Furthermore, as the threat landscape in the healthcare sector evolves and as new combative protection strategies emerge, it can be expected that government rules impacting healthcare information protection requirements will be updated. Readers responsible for healthcare sector security and privacy need to stay abreast of all such federal, state, or local government updates and make appropriate changes to their fielded security and privacy policies, procedures, and technologies as needed to meet new requirements.

71.7 PRIVATE SECTOR HEALTHCARE PRIVACY AND SECURITY INITIATIVES.

Various industry parties in the healthcare sector are engaging in industry initiatives and developing industry standards that pertain to healthcare information protection. The HITRUST efforts are examined in Section 71.7.1. Efforts focused on privacy and security as they apply to medical devices are described in Section 71.7.2.

71.7.1 HITRUST.

The Health Information Trust Alliance (HITRUST)[114] is an industry organization that creates industry standards for information security for various parties in the healthcare sector. HITRUST arose from an early initiative, the eHealth Vulnerability Reporting Program, that focused on finding, publicizing, and fixing vulnerabilities in healthcare application software, especially those software packages that were large and for which software development errors were more numerous, on average.

As presented below, HITRUST member organizations collaborated in developing a healthcare sector security framework, security specifications, and solutions for assessing and assuring compliance with claimed security implementations. HITRUST deliverables reflect best practices fielded in the healthcare industry.

HITRUST member organizations represent in-the-field private healthcare sector organizations that work directly with creating, managing, and exchanging electronic healthcare information as it flows through the healthcare sector.

71.7.1.1 HITRUST Security Initiatives. HITRUST looks to replace the existing hodgepodge of various commercial and/or proprietary silos of differing security and security assurance solutions currently used by employers, healthcare providers, and health insurance carriers. Instead, standard, comprehensive security, as well as security assessment practices, are specified for protecting, and assuring the protection of electronic healthcare information storage, access, and exchange.

The HITRUST Common Security Framework (CSF) is touted to provide for just this type of reasonable, flexible, trustworthy protection of healthcare information. Protection is provided by a common baseline of security capabilities plus, as appropriate, other additional, scaled, and/or emergent capabilities. The strategy is to make sure:

1. That a specific healthcare entity's security solution is appropriate and scaled to the specific healthcare entity's business and technology realities, to the specific risks to be managed, and to the specific healthcare entity's resources available to manage the risk, and
2. That the security solution can adapt to changing healthcare needs and evolving information and security technologies.

The CSF provides a harmonized approach for healthcare information security and privacy that reflects an amalgamation of protection concepts that have emerged across several parties, including federal laws and regulations (HIPAA, HITECH), federal government standards and guidelines (NIST, FTC, CMS), industry standards and regulations (ISO, PCI, COBIT), and state regulations. The CSF is available to HITRUST member organizations and to individuals from qualified organizations that do not provide security products or services.

71.7.1.2 HITRUST Compliance and Assurance Initiatives. HITRUST pursues a CSF compliance and assurance program,[115] According to HITRUST, the program is motivated and driven by:

1. The worsening threat landscape in the healthcare sector,
2. The cost of conducting various, multiple, often proprietary and inconsistent, conformance tests to satisfy different parties, and
3. The increasing penalties being levied on healthcare sector entities for noncompliance to various security and privacy specifications.

HITRUST compliance assurance is intended to elevate a healthcare entity's level of due diligence that appropriate security and privacy measures have been implemented in the conforming healthcare entity's environment. Of course, compliance to a comprehensive, appropriate, and adequate set of security mechanisms is key.

The HITRUST compliance and assurance program focuses on developing, and fostering healthcare sector use of HITRUST-standard assessment and reporting processes. These standard processes are used to determine and to document single-point-in-time compliance to any of several different federal and state security and privacy regulations, requirements and assurance laws, international standards, third-party standards,

healthcare entity business associate healthcare information protection requirements, CSF security controls, auditors' requirements, and more.

HITRUST offers different types of assessments to provide any of three different levels of assurance. For the first level of assurance, HITRUST makes available a Web-based tool to perform HITRUST-standardized self-assessments of various factors associated with a healthcare entity's security program. The factors include:

1. A healthcare entity's, or its business associates', security program's alignment with the compliance requirements associated with any variety of the several different security regulations and standards that may be applicable to the party whose conformance is being assessed,
2. The maturity of the security program, and
3. The program's ability to protect healthcare information in a sustainable way.

For use in higher risk environments, HITRUST supports two other levels of assurance—evaluated by either remote assessment or onsite assessment. All three HITRUST assessments are recognized across the healthcare sector.

The result of a compliance assessment is a report. HITRUST vets all conformance reports. Such reports include identification of any corrective actions needed to achieve compliance. CSF conformance assessment reports can be sent to various governing bodies for proof of compliance to required security and privacy measures.

To increase assurance/confidence that an HITRUST compliance assessment report is appropriate, accurate, and comprehensive, the assessment report can additionally be (a) validated by a periodically requalified third-party CSF assessor, as well as (b) certified by HITRUST.

71.7.1.3 Other HITRUST Initiatives. HITRUST runs a Cyberthreat Intelligence and Incident Coordination Center[116] and a Cyberthreat Analysis Service.[117] Together, these initiatives provide reports of emerging general threats and vulnerabilities, as well as threats and vulnerabilities specific to healthcare information protection. These initiatives also coordinate healthcare sector responses to potential cyberattacks on the healthcare sector. In addition, education is provided on pertinent topics such as healthcare technology risk analysis, healthcare business risk analysis, and more.

HITRUST efforts are also focused on developing healthcare sector-specific–guidelines, business partner contract templates, and protocols pertaining to physical, administrative, and technical security for healthcare sector information. Other HITRUST initiatives include providing healthcare information security education and awareness, analysis of evolving threats and solutions, commenting on and supporting the development of new federal rules applicable to healthcare information security, and more.

71.7.2 Medical Devices Security Initiatives. Medical devices are fixed function, electronic medical diagnosis, or treatment devices. Despite their limited function, such devices are susceptible to attack. Since medical devices often create, store, and convey PHI, they are under the jurisdiction of HIPAA and other information protection regulations. The risk of attack on medical devices will likely only escalate as more medical devices leave the confines of healthcare providers and become assigned to remote telemedicine workers or directly to patients.

Significant threats to PHI exposure from medical devices and to alteration of patient-specific treatment control signals executed by medical devices have been demonstrated

or speculated.[118] It has been shown that a wireless insulin pump that exchanges unencrypted healthcare information and pump controls with an external monitor or an external control system could be directed by a hacking intruder to be turned off or to dispense its entire insulin tank immediately. Defibrillators and pacemakers have allowed PHI to be compromised. They have also been remotely reprogrammed by intruders to permit denial-of-service attacks or to deliver hacker-directed high-voltage shocks. Networked medical devices that use commercial operating systems and are exposed to the Internet can be vulnerable to the same kinds of evolving vulnerabilities seen throughout all other industries. Lack of sufficient authentication and any encryption, use of software or firmware not tested for information protection, inability to patch software, use of unsupported legacy operating systems, and limited battery power, as well as limited CPU, memory, and storage to support the execution of security code, and more are all potentially suspect.

At the time of the writing of this chapter, the U.S. General Accountability Office examined what role the U.S. Food and Drug Administration should initiate to inspect fielded medical devices and to evaluate the security provided by new medical devices before they are released into the marketplace.

A trio of standards[119] from the International Electrotechnical Commission (IEC 62304) and the International Organization for Standardization (ISO 14971 and ISO 18348) specify a number of requirements pertaining to the development and testing of medical software, including application as well as security-related software within medical devices. IEC 62304 standardizes life cycle software requirements and software processes for medical device software development. ISO 14971 specifies a process for identifying medical device hazards and a companion risk-management process including, for example, creating and monitoring risk control mechanisms and evaluating control effectiveness. ISO 13485 specifies an organization's quality management needs for demonstrating the organization's ability to develop medical devices conformant to regulatory requirements such as HIPAA.

Also of value to the security and privacy accompanying medical devices, at the time of the writing of this chapter academic researchers were prototyping and investigating the feasibility of a firewall to observe communications to/from a medical device and to block external attacks on, or hijacking of, wireless, implantable or wearable, medical devices.[120]

71.8 SECURITY AND PRIVACY IMPACTS OF TECHNOLOGY TRENDS.

Like all other industries, the healthcare sector has embraced technology, in part to improve the efficacy of its care and business models. What appears to be emerging is that parts of the healthcare delivery model are evolving from a provider-centric model, where patients come to a care provider, to a patient-centric model, where providers electronically come to and interact with remote patients. New to this model is the emergence of patient-facing technologies that will generate, store, process, display, and communicate patient-specific data and treatment plans, as well as associated billing.

Due to the often conservative nature of many healthcare parties, emerging, cutting-edge technology may not always be rapidly adopted. Nonetheless, many leading-edge healthcare entities and/or individual employees and care providers are quickly adopting emerging IT infrastructure or medical device technologies that can impact the efficacy of healthcare.

But, the acceptance of new technology and its introduction to routine healthcare operations often increases the complexity of providing and managing protection. New

technologies will have their own threats and vulnerabilities. Furthermore, government regulations on providing healthcare sector information protection may likely move from today's point-in-time compliance demonstration toward a continuous compliance demonstration model. As such trends occur, healthcare entity security and privacy policies will need to be updated, more protection gaps will open in fielded security technologies, eventually fielded security technologies will become useless for new breeds of stealthier attackers and malware, and new strategies and implementations for protecting healthcare information will need to be developed and fielded.

There are many kinds of new technologies that are being introduced and that are starting to impact the security and privacy considerations and policies of the healthcare sector. Indeed, there is evidence that introduction of new technologies, together with any new health information breaches that arise from such new technologies, is overtaking regulation compliance as the main driver for organizations to provide information protection.[121] With rapid technology evolution and changes, organizations are beginning to recognize that information protection regulations may not be able to evolve fast enough to handle newer technologies and the new risks they may impose. Correspondingly, organizations are realizing that compliance to static or more slowly evolving regulations alone can be insufficient to provide information protection within their organizations.[122] Some feel that ongoing security testing is key to maintaining reasonable information protection in quickly evolving threat landscapes.[123]

All such technology introduction trends lead to companion trends in new forms of risks to, and attacks on, healthcare information. A common principle that appears to be emerging is that as healthcare information flows to and through ever newer IT technologies, the healthcare sector security model is moving more toward a data-centric security model wherein a healthcare entity's security perimeter is positioned at the data level rather than at a facility or network or server or IT service provider or mobile device level.

Examples of the kinds of new technologies that are emerging appear below.

71.8.1 Mobile and Personal Technology. With the proliferation of new patient-facing technologies, patients are taking more responsibility for their healthcare. Indeed, they are also offloading some of the administrative, triage, and care follow-up tasks previously under the responsibility of providers' staff. Such patient-side technologies are leading to a new style of digital relationship between patients, providers, and insurance carriers. New patient-facing technologies will support remote, personal, tele-health sessions, tele-medicine sessions, provider-patient consultations, provider-to-home-caregiver consultations, test results tracking, automatic invoicing and bill paying, and much more.

While potentially improving provider or patient experiences, use of nomadic, mobile personal devices such as tablet computers and smart phones are leading the way to a so-called post-PC world wherein the so-called Bring Your Own Device (BYOD) movement that relies on mobile and personal technologies presents new challenges to healthcare entities' security and privacy posture. The emergence of ever-newer mobile operating systems, mobile health applications, and communication mechanisms adds to security and privacy challenges. Social media use also presents security and privacy concerns.

State-of-the-art support of evolving, enhanced mobile and personal technologies appears certain as the healthcare sector will continue to evolve toward younger personnel steeped in use of such technologies.

71.8.2 Virtualized Technology. Software-as-a-service (SaaS), IT virtualization, and the outsourcing of healthcare applications and data storage to technology providers in any of the various different types of ever-more complex clouds are virtual technologies that may be adopted by healthcare entities. Such technologies may increase the overall complexity of a healthcare entity's IT environment and can present yet other security and privacy challenges. Such challenges, such as identity and access management issues, are often not unique to the healthcare sector. Challenges exist; for example, if no agreements are made about matters such as (a) what privacy, security, and security assurance accompany external or virtualized third-party software, host computers and/or services, (b) how compliance of the external third-party's security provisions to HITECH/HIPAA can be monitored or verified in real time by the contracting healthcare entity, and (c) how HIPAA/HITECH liabilities and penalties will be apportioned between the contracting healthcare entity and the outsourced entity. Indeed, some authorities argue that as of the time of the writing of this chapter, the cloud is not yet ready for mission-critical applications—healthcare or otherwise—in part because no universal security and security vetting standards yet exist across all private and public[124] healthcare environments. Nevertheless, vendors are beginning to offer HIPAA/HITECH-compliant cloud services and companion secure access from any device, including personal, mobile devices; and they are doing so as full-fledged business associates to healthcare entities.[125]

71.8.3 Personalized Medical Technology. With regard to healthcare information and healthcare information databases, considerations regarding new security and privacy requirements and policies will accompany the introduction of new personally wearable or implantable healthcare technology that creates and communicates PHI or is reactive to patient-specific external controls. Security, privacy, and individual data element masking considerations must also accompany the introduction of individual-patient-centric, personal health information databases.

71.8.4 Big Data Trends. Government authorities envision that significant benefit will arise from the HITECH push toward new healthcare system models. Patient medical information, as well as carrier payment data, across the country can be aggregated, as needed, into or from new, enormously sized, massively horizontal, deeply longitudinal, highly distributed, potentially replicated data stores of heterogeneous data that can be stored, organized, integrated, and used in new ways. The expectation is to analyze such data, coupled with research data, pharmaceutical data, bioinformatics data, and more, to improve care; to lower healthcare costs; to analyze disease trajectories; to hone diagnostic and predictive abilities; to identify emerging pharmaceutical and medical device opportunities; to identify best practices; to improve underperforming facilities, techniques, or providers; to enhance disease research, public health research, medical treatment research, direct patient care, and treatment outcome analysis; and to uncover many more opportunities that have yet to be envisioned.

For such mammoth healthcare information databases, data may reside in traditional database systems or an array of different unstructured data stores; the specific PHI data elements from a specific patient may be copied, moved, and/or repositioned in ad hoc ways; the data acquisition velocity across all patients is extraordinarily high; and the location of healthcare data can be very broadly distributed among a large array of loosely federated, potentially self-organizing, data stores. Traditional database and analysis approaches are being replaced by new database and analytic technologies

and methods for capturing, storing, managing, and orchestrating the processing and analysis of such complex, migrating, mammoth data—colloquially called "big data."

Privacy and security for big data in the healthcare sector is an ongoing challenge. There is no question that such a vast store of information provides a huge target and challenge to would-be "big score" data attackers. Traditional perimeter-oriented security and privacy may not apply. Establishing, sharing, and enforcing privacy consents (potentially among multiple copies of personal healthcare information), stifling ways of inferring personal identity from anonymized (e.g., partially redacted) PHI, and otherwise maintaining privacy in parallel and distributed healthcare information processing are examples of challenging privacy issues. Establishing security interoperability, data access knowledge, and end-to-end trust among a diverse array of differing, potentially self-organizing security technologies that may be sequentially used by differing parties with fluid boundaries to meet common security requirements espoused in government regulations is also nontrivial.

71.8.5 Other Factors Associated with Technology Trends. Large, potentially virtual, health information research databases will arise by using big data technologies to integrate and to aggregate the PHI of millions of patients nationwide. Appropriate methods of de-identifying personally identifiable healthcare information data elements need to accompany the introduction of big data technologies. Correspondingly, data mining technologies will need to be coupled with data element masking or other patient de-identification approaches when extracting healthcare knowledge from such massive aggregated healthcare information databases. Appropriate privacy and security measures need to accompany use of such technologies.

The opening of APIs and portals for accessing certain packages of health information from databases of aggregated health information will also have ramifications on the need for updated, flexible security and privacy policies and technologies. For example, in support of ONC goals to enhance patient engagement with electronic healthcare information, individual patients will likely rely on database portal technology to access their own information within the sea of healthcare data; but individual patients who have access to some specific small portion of healthcare data must not be allowed to access larger, cross-patient datasets that may be available to a public health researcher, to an insurer, or to appropriate government bodies for regulation refinement or law enforcement.

While the security and privacy challenges associated with emerging technology trends may not be insurmountable, healthcare sector parties do need to stay alert of emergent technology trends so that their security and privacy programs can be appropriately adjusted to accommodate technology trends being incorporated into the healthcare sector. Furthermore, while the emerging technologies used by the healthcare sector will change and potentially need new security and privacy mechanisms, the security and privacy capabilities in fielded security tools used by the healthcare community may also need to evolve. Over time, as common attack vectors evolve, or as one-of-a-kind, individually targeted attacks emerge, fielded security tools may no longer be able to keep healthcare information protected. New or updated security tools will need to be fielded.

However, healthcare sector parties need to stay aware of the possibility that even new fielded protection technologies and procedures may not completely be able to protect against the release of PHI. As security becomes more of a risk-management activity, it is envisioned that there can be a growing role for insurance in transferring or mitigating the risk and associated consequences and penalties of healthcare information theft

and leakage due to inadequate protection. The cost and amount of insurance coverage made available by an insurance carrier should be commensurate with the level of due diligence associated with a healthcare entity's fielded security policies and solutions.[126]

The bottom line is that creation, storage, communication, sharing, and use of healthcare information wherever it will exist will occur by means of an IT infrastructure with constantly moving parts operating under a growing number of government and industry compliance requirements. Healthcare sector parties will need to stay vigilant and nimble in terms of protecting their evolving technologies and processes and the PHI entrusted to them. Some may consider outsourcing their need for appropriately compliant security solutions to an MSSP (Managed Security Solutions Provider), that is, a business associate in the eyes of HIPAA/HITECH that specializes in providing compliant information protection services to the healthcare sector. However, the cost of maintaining information protection in the face of evolving technologies, risks, and regulations will compete with several other line items in a healthcare entity's budget. Corporate officers need to be properly engaged to allocate sufficient resources to keep up with evolving information protection needs.

71.9 SUMMARY. Appropriately protected personal healthcare information and healthcare business processes are at the heart of proper healthcare. But, there are innumerable and growing threats to their security and privacy.

Through regulations, the government is pushing the healthcare sector in two directions: (1) HIPAA and other government regulations are pushing to keep patient records safe, and (2) HITECH MU regulations are pushing to make patient records more accessible and broadly available.

The first direction leads to more frequently recurring, so-called "living," risk assessments by healthcare entities. Such continuous risk assessments will be driven as government steps up regulatory enforcement, as patient demand for information safety and protection increases, as private lawsuits against healthcare information custodians liable for personal information compromise catch on, and as contract law legal proceedings arise between covered entities and business associates.

The second direction leads toward migration to broad dissemination technologies, such as mobile, cloud, and regional health exchange technologies.

But, the combined impact of both directions is to increase vulnerabilities to data exposure and to expand the required breadth of protection for healthcare information. Increased vulnerabilities couple with evolving, more sophisticated, more diverse targeted threats and attacks to increase the amount of healthcare information being compromised.

There are multitudes of government and industry laws, regulations, and initiatives pertinent to protecting healthcare information. However, traditionally the healthcare sector has focused on acquiring ever-newer revenue-generating medical technologies. It has been slow to turn its corporate attention and funds toward information protection, often doing just the minimum in terms of complying with government laws and regulations.

Ironically, looking across the healthcare sector, compliance with government information protection regulations has not seemed to have had a profound impact on personal healthcare information privacy protection: Healthcare information breaches are rampant despite adherence to government security regulations. The focus on compliance has not always led to better healthcare information protection.

Patient care typically trumps healthcare information protection, both in terms of capital and labor expenditures. In fact, the fraction of a typical healthcare organization

budget allocated to information protection tends to be half, if not less, than that seen in other industries.[127] Even government compliance and breach enforcement fines on healthcare entities do not necessarily prevent repeat offenses of healthcare information compromise.

Privacy and security leaders serving healthcare entities must be able to express healthcare information protection needs in ways to gain healthcare entity executive support. For example, they must show executives how information protection critically supports the primary, moneymaking, care delivery business of the healthcare entity.

The current government drive for healthcare improvements while forcing reductions in healthcare costs and insurance carrier reimbursements does not help the government-mandated healthcare information protection situation. Healthcare organizations and professionals are caught in a bind and are under pressure to work out their corporate and Hippocratic priorities and regulatory responsibilities.

71.10 FURTHER READING

Anderson, R. *Personal Medical Information: Security, Engineering, and Ethics*. Springer, 1997 (reprinted 2008).

Casscells, S. W. *DoD Health Information Security Regulation*. U.S. Department of Defense (July 12, 2007).

Davis, N. A., and M. LaCour. *Health Information Technology*, 3rd ed. Saunders, 2013.

Jones, E. D., and C. P. Hartley. *Policies and Procedures for the Electronically Connected Medical Office*. American Medical Association, 2009.

Kabay, M. E. "Medical Informatics Security." (May 2004). www.mekabay.com/overviews/medical.pdf

Nelson, R., and N. Staggers. *Health Informatics: An Interprofessional Approach*. Mosby, 2013.

Shoniregun, C. A., K. Dube, and F. Mtenzi. *Electronic Healthcare Information Security*. Springer, 2010 (reprinted 2013).

Trotter, F., and D. Uhlman. *Hacking Healthcare: A Guide to Standards, Workflows, and Meaningful Use*. O'Reilly Media, 2011.

Wager, K. A., F. W. Lee, and J. P. Glaser. *Health Care Information Systems: A Practical Approach for Health Care Management*, 3rd ed. Jossey-Bass, 2013.

71.11 NOTES

1. "ACO Participants Investing in Analytics, Survey Finds," *Healthcare IT News | HIEWatch*, March 15, 2013, www.hiewatch.com/news/aco-participants-investing-analytics-survey-finds
2. A more extensive listing of many more of the federal laws and regulations addressing privacy and security in the healthcare sector together with brief summaries is available at www.healthit.gov/sites/default/files/federal_privacy_laws_table_2_26_10_final_0.pdf
3. "FireEye Advanced Threat Report 2 H 2012," http://www2.fireeye.com/rs/fireye/images/fireeye-advanced-threat-report-2h2012.pdf
4. R. Rachwald, "The New FireEye Advanced Threat Report," FireEye Blog, Advanced Malware, Threat Intelligence, April 3, 2013, www.fireeye.com/blog/technical/malware-research/2013/04/the-new-fireeye-advanced-threat-report.html

5. Thor Olavsrud, "Healthcare Industry CIOs, CSOs Must Improve Security," *Network World* newsletter, March 6, 2012.
6. Nicole Lewis, "Health Data Breaches Up 97% in 2011," *InformationWeek | Healthcare*, February 13, 2012, www.informationweek.com/healthcare/security-privacy/health-data-breaches-up-97-in-2011/232600746?cid=nl_IW_healthcare_2012-02-15_html&elq=022382be2aa443e9a3b3fd4dba6ce46f
7. Identity Theft Resource Center, "2012 Data Breach Stats," December 26, 2012, www.idtheftcenter.org/images/breach/Breach_Stats_Report_2012.pdf
8. M. Millard, "Breaches epidemic despite efforts at compliance, says Kroll," *Healthcare IT News*, April 10, 2012, www.healthcareitnews.com/news/breaches-epidemic-despite-efforts-compliance-says-kroll
9. Kelly Jackson Higgins, "Most Small Healthcare Practices Hacked in the Past 12 Months," *Dark Reading | Database | News*, February 16, 2012, www.darkreading.com/database-security/167901020/security/news/232601045/most-small-healthcare-practices-hacked-in-the-past-12-months.html?cid=nl_DR_daily_2012-02-17_html&elq=eeb7dbbf00374e629c9db74fae6ebf54
10. Nicole Lewis, "Patient Data Theft Sends IT Specialist to Jail," *InformationWeek | Healthcare*, January 17, 2012, www.informationweek.com/healthcare/security-privacy/patient-data-theft-sends-it-specialist-t/232400459?cid=nl_IW_healthcare_2012-01-18_html&elq=ab0ba4c34056469798ce4eb49d3fb11 d
11. D. Walker, "Medical Identity Theft To Be Explored at FTC Hearing," *SC Magazine*, May 3, 2013, www.scmagazine.com/medical-identity-theft-to-be-explored-at-ftc-hearing/article/291780/?DCMP=EMC-SCUS_Newswire
12. RSA, The Security Division of EMC^2, "Cybercrime and the Healthcare Industry," White Paper, 2010, www.rsa.com/products/consumer/whitepapers/11030_CYBHC_WP_0710.pdf (URL inactive).
13. Ericka Chickowski, "Healthcare Unable to Keep Up With Insider Threats," *Dark Reading | Insider Threat*, | *News*, May 1, 2012, www.darkreading.com/insider-threat/healthcare-unable-to-keep-up-with-inside/232901235
14. U. S. Department of Health and Human Services, "Health Information Privacy: Breaches Affecting 500 or More Individuals," HHS Website, www.hhs.gov/ocr/privacy/hipaa/administrative/breachnotificationrule/breachtool.html
15. J. Taitsman, C. M. Grimm, and S. Agrawal, "Protecting Patient Privacy and Data Security," *New England Journal of Medicine* 368 (2013):977–979, March 14, 2013, www.nejm.org/doi/full/10.1056/NEJMp1215258
16. Mathew J. Schwartz, "Hackers Holds Australian Medical Records Ransom," *InformationWeek | Security*, December 11, 2012, www.informationweek.com/security/attacks/hackers-hold-australian-medical-records/240144164?cid=nl_IW_healthcare_2012-12-12_html&elq=f71d081ae23a4929bdc01e6b0c9b37f1
17. T. Armerding, "Law Firms See Big Money in Healthcare Breach Cases," *CSO Online | Data Protection*. April 16, 2012, www.csoonline.com/article/704288/law-firms-see-big-money-in-healthcare-breach-cases?page=1
18. A. Pascual, "Financial Pain Ensues When Custodians of Health Fail to be Good Stewards of Privacy," Javelin Strategy & Research Blog, April 28, 2013, https://www.javelinstrategy.com/blog/2013/04/28/financial-pain-ensues-when-custodians-of-health-fail-to-be-good-stewards-of-privacy (see also E. McCann, "Link Between Breaches and Fraud on the Rise," *Healthcare IT News*

| *HIEWatch*, April 30, 2013, www.hiewatch.com/news/link-between-breaches-and-fraud-rise)

19. Ericka Chickowski, "$1.5M Fine Marks a New Era in HITECH Enforcement," *Dark Reading | Database | News*, March 21, 2012, www.darkreading.com/database/15 m-fine-marks-a-new-era-in-hitech-enfor/232700031
20. J. Robbins, "Cybertheft Heists $1 Million from Leavenworth Hospital," *The Wenatchee World*, April 26, 2013, www.wenatcheeworld.com/news/2013/apr/26/cybertheft-heists-1-million-from-leavenworth
21. Ken Terry, "Should FDA Assess Medical Device Defenses against Hackers?" *InformationWeek | Healthcare*, April 24, 2012, www.informationweek.com/healthcare/security-privacy/should-fda-assess-medical-device-defense/232900818?cid=nl_IW_healthcare_2012-04-25_html&elq=b34ee79cddae4a828974ed6533178f9 a
22. Kim Zetter, "Vulnerability Lets Hackers Control Building Locks, Electricity, Elevators and More," *Wired | Threat Level*, February 6, 2013, www.wired.com/threatlevel/2013/02/tridium-niagara-zero-day
23. Kelly Jackson Higgins, "Medical Industry Under Attack By Chinese Hackers," *Dark Reading | Threat Intel | News*, March 14, 2013, www.darkreading.com/threat-intelligence/167901121/security/attacks-breaches/240150858/medical-industry-under-attack-by-chinese-hackers.html?cid=nl_DR_daily_2013-03-15_html&elq=da19ca4090704bcd87b093ad6ea14710
24. Cyber^2, "Medical Industry—A Cyber Victim: Billions Stolen and Lives at Risk," White Paper, March 14, 2013, www.cybersquared.com/wp-content/uploads/downloads/2013/03/Medical-Industry-A-Cyber-Victim-Billions-Stolen-and-Lives-At-Risk.pdf
25. E. Chickowski, "Healthcare Securing Improving But Still Needs Treatment," *Dark Reading | Database Security | News*, March 27, 2013, www.darkreading.com/database-security/167901020/security/news/240151816/in-spite-of-improvement-healthcare-security-still-needs-treatment.html?cid=nl_DR_daily_2013-03-27_html&elq=7a12116be5bc4b20b2b318a4ebf967 dd
26. Solutionary, "2013 Global Threat Intelligence Report (GTIR)."
27. Kevin Fogarty, "DNA Hack Could Make Medical Privacy Impossible," *NetworkWorld | Security*, March 11, 2013, www.networkworld.com/news/2013/031113-dna-hack-could-make-medical-267569.html?source=NWWNLE_nlt_daily_pm_2013-03-11
28. National Conference of State Legislatures, "Security Breach Notification Laws," NCSL Website, August 20, 2012, www.ncsl.org/issues-research/telecom/security-breach-notification-laws.aspx
29. 104th Congress of the USA, "Health Insurance Portability and Accountability Act of 1996," Public Law 104-191, Code of Federal Regulations (CFR) Title 45 Part 164 (August 21, 1996), http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=104_cong_public_laws&docid=f:publ191.104.pdf or http://tinyurl.com/5z4 csd
30. U. S. Department of Health and Human Services, OCR, "HIPAA Administrative Simplification Statute and Rules," HHS.gov Website, http://www.hhs.gov/ocr/privacy/hipaa/administrative
31. HIPAA security requirements for specified types of healthcare transactions were described in Chapter 71 of the 5th edition of the *Computer Security Handbook*.

32. CMS.gov, Centers for Medicare & Medicaid Services, "Are You a Covered Entity?" www.cms.gov/Regulations-and-Guidance/HIPAA-Administrative-Simplification/HIPAAGenInfo/AreYouaCoveredEntity.html, and "Covered Entity Charts," www.cms.gov/Regulations-and-Guidance/HIPAA-Administrative-Simplification/HIPAAGenInfo/Downloads/CoveredEntitycharts.pdf
33. Analog forms of information such as paper-to-paper fax, voice telephony, mail, and so on are not within HIPAA's scope.
34. U.S. Department of Health and Human Services, Office of the Secretary, "Standards for Privacy of Individually Identifiable Health Information," 45 CFR Parts 160 and 164, *Federal Register* 65, no. 250 (December 28, 2000), www.hhs.gov/ocr/privacy/hipaa/administrative/privacyrule/prdecember2000all8parts.pdf
35. U.S. Department of Health and Human Services, Office of the Secretary, "Standards for Privacy of Individually Identifiable Health Information," 45 CFR Parts 160 and 164, *Federal Register* 67, no. 157 (August 14, 2002), www.hhs.gov/ocr/privacy/hipaa/administrative/privacyrule/privruletxt.txt
36. H. Hayes, "HIPAA: Best if Used By," *Government Health IT*, June 12, 2006, www.govhealthit.com/news/hipaa-best-if-used
37. U.S. Department of Health and Human Services, Office of the Secretary, "Health Insurance Reform: Security Standards," 45 CFR Parts 160, 162, and 164, *Federal Register* 68, no. 34 (February 20, 2003), www.hhs.gov/ocr/privacy/hipaa/administrative/securityrule/securityrulepdf.pdf
38. M. Scholl, K. Stine, J. Hash, P. Bowen, A. Johnson, C. D. Smith, and D. I. Steinberg, "An Introductory Resource Guide for Implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule," NIST Spec Pub 800-66 Rev1, October 2008, www.hhs.gov/ocr/privacy/hipaa/administrative/securityrule/nist80066.pdf
39. U.S. Department of Health and Human Services, OCR, "Breach Notification for Unsecured Protected Health Information," 45 CFR Parts 160 and 164, *Federal Registry* 74, no. 162 (August 24, 2009): 42740–42770, www.gpo.gov/fdsys/pkg/FR-2009-08-24/pdf/E9-20169.pdf
40. Federal Trade Commission, "Health Breach Notification Rule; Final Rule," 16 CFR Part 318, *Federal Register* 74, no. 163 (August 25, 2009), http://www2.ftc.gov/os/2009/08/R911002hbn.pdf
41. U.S. Department of Health and Human Services, Office of the Secretary, "Guidance Specifying Technologies and Methodologies That Render Protected Health Information Unusable, Unreadable or Indecipherable to Unauthorized Individuals for Purposes of the Breach Notification Requirements," 45 CFR Parts 160 and 164, *Federal Register* 74, no. 79 (April 27, 2009), www.hhs.gov/ocr/privacy/hipaa/understanding/coveredentities/federalregisterbreachrfi.pdf
42. It is possible that breach detection may arise as a consequence of a formal complaint of a HIPAA privacy or security rules violation filed with the HHS OCR.
43. U.S. Department of Health and Human Services, "How to File a Complaint," HHS.gov | Health Information Privacy, www.hhs.gov/ocr/privacy/hipaa/complaints/index.html
44. U.S. Department of Health and Human Services, "What OCR Considers During Intake & Review," HHS.gov | Health Information Privacy, www.hhs.gov/ocr/privacy/hipaa/enforcement/process/whatocrconsiders.html

45. U.S. Department of Health and Human Services, "How OCR Enforces the HIPAA Privacy Rule," HHS.gov | Health Information Privacy, www.hhs.gov/ocr/privacy/hipaa/enforcement/process/howocrenforces.html
46. U.S. Department of Health and Human Services, "Modifications to the HIPAA Privacy, Security, Enforcement, and Breach Notification Rules Under the Health Information Technology for Economic and Clinical Health Act and the Genetic Information Nondiscrimination Act; Other Modifications to the HIPAA Rules," *Federal Register*, Document Citation 78 FR 5565, Doc No. 2013-01073 (January 24, 2013): 5565–5702, https://federalregister.gov/a/2013-01073
47. Goodwin Procter LLP, "HHS Issues Long-Awaited Final HIPAA Omnibus Rule," *Legal Alert Newsletter*, January 30, 2013.
48. U.S. Department of Health and Human Services, Office of the Secretary, "Modifications to the HIPAA Privacy, Security, Enforcement, and Breach Notification Rules Under the Health Information Technology for Economic and Clinical Health Act and the Genetic Information Nondiscrimination Act; Other Modifications to the HIPAA Rules, Final Rule," 45 CFR 160 and 164, *Federal Register*, Document Citation 78 FR 5565, Doc No. 2013-01073 (January 25, 2013): 5565–5702, https://federalregister.gov/a/2013-01073
49. FTC Breach Notification Rule.
50. A. Smith, "HIPAA Rule Alters Definition of 'Breach,' " Society for Human Resource Management | Legal Issues | Federal Resources, January 28, 2013, www.shrm.org/LegalIssues/FederalResources/Pages/HIPAA-rule-breach.aspx
51. Legal & Regulatory Affairs Staff, "HIPAA Final Rule Highlights for Practitioners," Practice Central, March 14, 2013, www.apapracticecentral.org/update/2013/03-14/final-rule.aspx
52. P. Brusil, "Medical Records Protection," Chapter 71, Section 71.6.6.3, in *Computer Security Handbook*, 5th ed., ed. S. Bosworth, M. E. Kabay, and E. Whyne, Hoboken, NJ: John Wiley & Sons, Inc., 2009.
53. Healthcare entities often have to demonstrate compliance to several other regulations. Bodies associated with the different regulations may provide guidance on how best to meet such other compliance requirements. For example, the PCI Security Standards Council publishes such guidelines for the various updates to the PCI DSS. As needed, interested readers should examine such other documents.
54. HIPAA Certification, "HIPAA Compliance Certification: Covered Entity," www.hipaacertification.net/HIPAA-compliance-certification.htm
55. Some observers argue that going through a spot-check OCR audit or enduring the consequences of an OCR or Attorney General investigation into a reported breach has not always been shown to alter the information protection behavior of the impacted healthcare entity.
56. J. Roman, "HIPAA Audits: The Next Phase; OCR's Leon Rodriguez Offers Compliance Tips," Healthcare Info Security, January 2, 2013, www.healthcareinfosecurity.com/hipaa-audits-next-phase-a-5397
57. ENHAC, "ENHAC to Fill HIPAA, Meaningful Use Compliance Gaps," www.ehnac.org/press-news/press-articles/194-hopkinson (URL inactive).
58. URAC, Healthcare Organizations, "Welcome to the Health IT Resource Center," https://www.urac.org/infoTech (see also "HIPAA Privacy Accreditation Program Overview," https://www.urac.org/programs/prog_accred_HIPAAP_po.aspx?

navid=accreditation&pagename=prog_accred_HIPAAP and "HIPAA Security Accreditation Program Overview," https://www.urac.org/programs/prog_accred_HIPAAS_po.aspx)

59. "Health IT Privacy and Security Resources," HealthIT.gov | Providers & Professionals | Privacy & Security, www.healthit.gov/providers-professionals/ehr-privacy-security/resources
60. "Health Information Privacy and Security: A 10 Step Plan," HealthIT.gov | Providers & Professionals | Privacy & Security, www.healthit.gov/providers-professionals/ehr-privacy-security/10-step-plan
61. U.S. Department of Health and Human Services, "[Final] Guidance on Risk Analysis Requirements under the HIPAA Security Rule," Office for Civil Rights, www.hhs.gov/ocr/privacy/hipaa/administrative/securityrule/rafinalguidancepdf.pdf
62. U.S. Department of Health and Human Services, "Health Information Technology," HIPAA Privacy Components of the ONC Privacy and Security Toolkit, Office for Civil Rights | Health Information Privacy, www.hhs.gov/ocr/privacy/hipaa/understanding/special/healthit
63. U.S. Department of Health and Human Services, "Security Rule Guidance Material," Office for Civil Rights | Health Information Privacy, www.hhs.gov/ocr/privacy/hipaa/administrative/securityrule/securityruleguidance.html
64. U.S. Department of Health and Human Services, "HIPAA Security Guidance," Office for Civil Rights, www.hhs.gov/ocr/privacy/hipaa/administrative/securityrule/remoteuse.pdf
65. U.S. Department of Health and Human Services, "Health IT Privacy and Security Resources," HealthIT.gov | Privacy & Security, www.healthit.gov/providers-professionals/ehr-privacy-security/resources
66. U.S. Department of Health and Human Services, "Your Mobile Device and Health Information Privacy and Security," HealthIT.gov | Privacy & Security, www.healthit.gov/providers-professionals/your-mobile-device-and-health-information-privacy-and-security
67. U.S. Department of Health and Human Services, Security Rule Educational Paper Series, HHS.gov | Health Information Privacy, www.hhs.gov/ocr/privacy/hipaa/administrative/securityrule/securityruleguidance.html
68. National Institute of Standards and Technology, " HIPAA Security Rule Toolkit," NIST Information Technology Laboratory | Security Content Automation Protocol, http://scap.nist.gov/hipaa
69. M. Scholl et al., "An Introductory Resource for Implementing the Health Insurance Portability and Accountability (HIPAA) Security Rule."
70. U.S. Department of Health and Human Services, "How Can You Protect and Secure Health Information When Using a Mobile Device?" HealthIT.gov | Providers & Professionals | Mobile Device Privacy and Security, www.healthit.gov/providers-professionals/how-can-you-protect-and-secure-health-information-when-using-mobile-device
71. U.S. Department of Health and Human Services, "Your Mobile Device and Health Information Privacy and Security," HealthIT.gov | Providers & Professionals, | Privacy & Security, www.healthit.gov/providers-professionals/your-mobile-device-and-health-information-privacy-and-security

72. M. Souppaya and K. Scarfone, "Guidelines for Managing and Securing Mobile Devices in The Enterprise," National Institute of Standards and Technology, NIST Special Publication 800-124, July 2012 (or succeeding revisions thereof).
73. L. Chen, J. Franklin, and A. Regenscheid, "Guidelines on Hardware-Rooted Security in Mobile Devices," National Institute of Standards and Technology, NIST Special Publication 800-164, October 2012 (or succeeding revisions thereof).
74. U.S. Department of Health and Human Services, "Other Security Rule Notices and Materials," Office for Civil Rights | Health Information Privacy, www.hhs.gov/ocr/privacy/hipaa/administrative/securityrule/index.html
75. American Hospital Association, "HIPAA State Preemption Analysis," HIPAA Privacy Standards, www.aha.org/content/00-10/preemption.pdf
76. A. Brino, "When HIPAA and States Collide," Healthcare IT News | HIE Watch, April 18, 2013, www.hiewatch.com/perspective/when-hipaa-and-states-collide
77. U.S. Department of Health and Human Services, "Health Information Security & Privacy Collaboration (HISPC)," HealthIT.gov/Policymaking, Regulation, & Strategy | Federal-State Health care Coordination, www.healthit.gov/policy-researchers-implementers/health-information-security-privacy-collaboration-hispc
78. U.S. Department of Health and Human Services, "Health Information Security & Privacy Collaboration (HISPC)."
79. George W. Bush, "Incentives for the Use of Health Information Technology & Establishing the Position of the NHI IT Coordinator," Executive Order EO 13335, April 27, 2004. Follow-on to Bush, "Transforming Health Care: The President's Health Information Technology Plan," http://georgewbush-whitehouse.archives.gov/infocus/technology/economic_policy200404/chap3.html
80. Health Information Technology for Economic and Clinical Health Act, Title XIII of the American Recovery and Reinvestmant Act of 2009, Public Law 111-5.
81. The MU initiative was created by the American Recovery Act of 2008. The ONC administers and leads the MU initiative. The MU program incorporates government-paid incentive payments through 2015 to coax providers to adopt technologies and processes that meet MU requirements and goals.
82. Such groups include, for example, the ONC Information Exchange Working Group (IEWG), the Massachusetts eHealth Collaborative, the electronic Health Record Association (EHRA), and the CommonWealth Health Alliance.
83. Ken Terry, "EHRA Guides Vendors on Data Transport Mechanism," *InformationWeek* | Healthcare, March 5, 2013, www.informationweek.com/healthcare/electronic-medical-records/ehra-guides-vendors-on-data-transport-me/240149830?cid=nl_IW_healthcare_2013-03-06_html&elq=bac8e913933c4b41952210a6b85116b0
84. U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology, Health Information Technology Policy Committee, "Request for Comment Regarding the Stage 3 Definition of Meaningful Use of Electronic Health Records (EHRs)," for comments to be delivered by January 14, 2013, www.healthit.gov/sites/default/files/hitpc_stage3_rfc_final.pdf

85. S&I Framework, "Data Segmentation for Privacy Charter and Members," S&I Framework Wiki, http://wiki.siframework.org/Data+Segmentation+for+Privacy+Charter+and+Members
86. "Roberta," "Electronically Exchanging Sensitive Health Information", HIE Answers Website, September 19, 2012, www.hieanswers.net/electronically-exchanging-sensitive-health-information
87. J. Coleman and D. Decouteau, ONC, "Use of Interoperability Standards and Data Segmentation to Support Patient Privacy," presentation to HIMSS13 Conference, New Orleans, LA, 2013, http://wiki.siframework.org/file/view/DS4P%20Showcase_v3FINAL.pdf/418765700/DS4P%20Showcase_v3FINAL.pdf
88. U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology, "Nationwide Privacy and Security Framework for Electronic Exchange of Individually Identifiable Health Information," December 15, 2008, www.healthit.gov/sites/default/files/nationwide-ps-framework-5.pdf
89. U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology, "Guide to Privacy and Security of Health Information," Version 1.2 060112, www.healthit.gov/sites/default/files/pdf/privacy/privacy-and-security-guide.pdf
90. U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology, "Guide to Privacy and Security of Health Information," Chapter 2 "Privacy & Security and Meaningful Use," www.healthit.gov/sites/default/files/pdf/privacy/privacy-and-security-guide-chapter-2.pdf
91. U.S. Department of Health and Human Services, Office of the Secretary, "Health Information Technology: Standards, Implementation Specifications, and Certification Criteria for Electronic Health Record Technology, 2014 Edition; Revisions to the Permanent Certification Program for Health Information Technology" (Final Rule), *Federal Register* 77, no. 171 (September 4, 2012), www.gpo.gov/fdsys/pkg/FR-2012-09-04/pdf/2012-20982.pdf
92. U.S. Department of Health and Human Services, Office of the National Coordinator for Health Information Technology, "Certified Health IT Product List," HealthIT.gov Website, http://oncchpl.force.com/ehrcert
93. Critics of the Meaningful Use program point out that "attesting" to MU requirements pertaining to HIPAA is not as strong a show of compliance to HIPAA than,for example, demonstrating Meaningful Use of privacy and security capabilities that meet HIPAA requirements. Furthermore, attestation can lead to fraud if providers falsely attest to MU capabilities and inappropriately obtain MU incentive payments from the Centers for Medicare and Medicaid Services.
94. U.S. Department of Health and Human Services, "How Do I Comply with Meaningful Use Requirements," Health Resources and Service Administration Website, www.hrsa.gov/healthit/toolbox/HIVAIDSCaretoolbox/SecurityAndPrivacyIssues/howdoicomplywitmu.html
95. Centers for Medicare & Medicaid Services.
96. "The Direct Project," http://wiki.directproject.org/home

97. U.S. Department of Health and Human Services, "State Health Information Exchange," HealthIT.gov | Providers & Professionals | Help Center, www.healthit.gov/providers-professionals/state-health-information-exchange
98. B. Andriesen, "Bright Spots of the Direct Project Adoption," HealthIT.gov | HealthIT Buzz blog, June 18, 2012, www.healthit.gov/buzz-blog/state-hie/the-direct-project-adoption
99. K. Terry, "Healthcare Directed Exchange Gets Big Government Push," *InformationWeek* | Healthcare | News, April 16, 2013, www.informationweek.com/healthcare/interoperability/healthcare-directed-exchange-gets-big-go/240152932?cid=nl_IW_healthcare_2013-04-17_html&elq=1bfc583e686e4c96a60149885848e530
100. W. Burr, D. Dodson, E. Newton, R. Perlner, T. Polk, S. Gupta, and E. Nabbus, "Electronic Authentication Guideline," NIST Special Publication 800-63-1, December 2011, http://csrc.nist.gov/publications/nistpubs/800-63-1/SP-800-63-1.pdf
101. Federal Information Security Management Act of 2002, Title III of the E-Government Act of 2002, Pub.Law 107-347, December 17, 2002, http://csrc.nist.gov/drivers/documents/FISMA-final.pdf
102. Joint Task Force Transformative Initiative, "Security and Privacy Controls for Federal Information Systems and Organizations," NIST Special Publication 800-53 Rev 4, Final Public Draft (to be superseded by the final version), February 5, 2013, http://csrc.nist.gov/publications/drafts/800-53-rev4/sp800_53_r4_draft_fpd.pdf
103. M. E. Kabay, "NIST SP800-53 Rev. 3: Key to Unified Security Across Federal Government and Private Sectors" (Part 1), *NetworkWorld* | Security Strategies Alert, October 26, 2009, www.networkworld.com/newsletters/sec/2009/102609sec1.html
104. M. E. Kabay, "NIST SP800-53 Rev. 3: Risk Management Framework Underpins the Security Life Cycle" (Part 2), *NetworkWorld* | Security Strategies Alert, October 28, 2009, www.networkworld.com/newsletters/sec/2009/102609sec2.html
105. M. E. Kabay, "Extensive Catalog Provides Security Controls for Contemporary Security Requirements" (Part 3), *NetworkWorld* | Security Strategies Alert, November 2, 2009, www.networkworld.com/newsletters/sec/2009/110209sec1.html
106. M. E. Kabay, "SP 800-53 is Essential for Security in Federal Government IT Systems" (Part 4), *NetworkWorld* | Security Strategies Alert, November 4, 2009, www.networkworld.com/news/2009/110209-applicability-government.html
107. Many NIST SP 800-53 security controls map to security controls defined in international security standard ISO/IEC 27001.
108. FedRAMP, "Ensuring Secure Cloud Computing for the Federal Government," U.S. General Services Administration Website, www.gsa.gov/portal/category/102371
109. See, for example, www.niap-ccevs.org/CCEVS_Products/pcl.cfm or www.commoncriteriaportal.org/products
110. Kim Zetter, "Board Urges Feds to Prevent Medical Device Hacking," *Wired* | Threat Level: Cybersecurity, April 10, 2012, www.wired.com/threatlevel/2012/04/security-of-medical-devices

111. Such a program was attempted by NIST in the late 1990s, whereby healthcare security requirements would be specified by Common Criteria methods; government-sanctioned, third-party, security testing labs would evaluate products to verify that the stipulated security requirements were met; and the government would validate the product evaluations and provide a product evaluation certificate.
112. Barack Obama, "Executive Order 13636—Improving Critical Infrastructure Cybersecurity," February 12, 2013, www.whitehouse.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity
113. Barack Obama, "Critical Infrastructure Security and Resilience," Presidential Policy Directive PDD-21, February 12, 2013, www.fas.org/irp/offdocs/ppd/ppd-21.pdf
114. Information about HITRUST and its programs and deliverables is available at the HITRUST Website, www.hitrustalliance.net
115. HITRUST, "HITRUST CSF Assurance Program," Brochure, www.hitrustalliance.net/HIT_Bro_CSF_Assure_v2_Feb_2013.pdf
116. HITRUST, "HITRUST Cyber Threat Intelligence and Incident Coordination Center," www.hitrustalliance.net/c3
117. HITRUST, "HITRUST Cyber Threat Analysis Service," www.hitrustalliance.net/ctas
118. D. Bradbury, "The Prognosis for Medical Device Security," Infosecurity Magazine Q2 (2013):41–43, http://content.yudu.com/Library/A24a1q/InfoSecMagQ2V2/resources/42.htm
119. American National Standards Institute (ANSI), "ANSI/AAMI/ISO 13485, ANSI/AAMI/IEC 62304 and ANSI/AAMI/ISO 14971—Software, Quality and Risk Management of Medical Devices," Medical Devices Package, ANSI eStandards Store, http://webstore.ansi.org/RecordDetail.aspx?sku=ANSI%2fAAMI%2fISO+13485+%2f+ANSI%2fAAMI%2fIEC+62304+%2f+ANSI%2fAAMI%2fISO+14971+-+Medical+Devices+Package#.UUmyu6VXL8 t
120. United Press International, "Firewall Can Stop Medical Device 'Hacking,' " UPI.com | Science News, April 12, 2012, www.upi.com/Science_News/2012/04/12/Firewall-can-stop-medical-device-hacking/UPI-19931334276290
121. Vormetric, "New Technologies Surpass Compliance as Driver for Organizations to Secure Data," ITBusinessEdge, May 8, 2013, www.itbusinessedge.com/slideshows/new-technologies-surpass-compliance-as-driver-for-organizations-to-secure-data.html
122. With the advent of FISMA, it was first realized that regulated organizations were focusing more on regulatory compliance and the fear of regulatory audit rather than on security and actual risk reduction. Apparently, executives in affected organizations found compliance easier to measure. They also had penalties assessed based on lack of compliance not on appropriate risk management in an environment of advancing technology and complexity changes.
123. Spirent, "Achieving Cyber Security Readiness Within an Evolving Threat Landscape," February 13, 2013, in SC Magazine | White Paper Roundup, May 17, 2013, http://forms.madisonlogic.com/FormConfirmation.aspx?pub=18&pgr=259&src=2318&cmp=5303*=25776&frm;=446&pd=161275-9-6886-57-321-0 (URL inactive).

124. FedRAMP provides security assessment standards only for cloud software providers serving government clients.
125. T. Olavsrud, "Box Takes Electronic Health Records to the Cloud", CIO.com | Applications, April 25, 2013, www.cio.com/article/732375/Box_Takes_Electronic_Health_Records_to_the_Cloud?page=1&taxonomyId=600000
126. Bill Brenner, "Video: David Navetta on Cyber Insurance, Risk Management," CSO | Security and Risk | Salted Hash blog, March 15, 2013, http://blogs.csoonline.com/security-industry/2620/video-david-navetta-cyber-insurance-risk-management?source=CSONLE_nlt_update_2013-03-17
127. Karen Epper Hoffman, "Urgent Care—Health Providers Have Pressing Reasons to now Embrace Security," SC Magazine, April 1, 2013, www.scmagazine.com/urgent-care-safeguarding-data-at-health-care-providers/article/284420

CHAPTER 72

LEGAL AND POLICY ISSUES OF CENSORSHIP AND CONTENT FILTERING

Lee Tien, Seth Finkelstein, and Steven Lovaas

72.1 INTRODUCTION 72·1

72.1.1 Scope of This Chapter: Government Intervention 72·2

72.1.2 Whose Laws? Whose Standards? 72·2

72.1.3 Defining Objectionable Material: International Differences 72·3

72.2 U.S. CONTEXT: FIRST AMENDMENT RIGHTS 72·12

72.2.1 What Does the First Amendment Protect? 72·12

72.2.2 Basic First Amendment Principles 72·12

72.2.3 Limitations on Government Interference with Speech 72·14

72.2.4 Exceptions Where Speech Can Legally Be Limited 72·16

72.2.5 Legislation and Legislative Initiatives in the United States 72·18

72.2.6 Attempts to Control Access: Case Law 72·19

72.3 PARENTAL INVOLVEMENT/ RESPONSIBILITY 72·21

72.4 SUMMARY 72·22

72.5 FURTHER READING 72·23

72.6 NOTES 72·23

72.1 INTRODUCTION

Everyone has the right to freedom of opinion and expression; this right includes freedom to hold opinions without interference and to seek, receive, and impart information and ideas through any media and regardless of frontiers.[1]

One might think that the Internet will make this ringing proclamation a reality. Like no other technology, the Internet transcends national borders and eliminates barriers to the free flow of information. Governments, however, are controlling speech on the Internet.

72.1.1 Scope of This Chapter: Government Intervention. Many nations protect rights of free speech and free expression. The *First Amendment to the U.S. Constitution* provides that "Congress shall make no law … abridging the freedom of speech or of the press." The *Canadian Charter of Rights and Freedoms* protects freedom of speech and press in Article 2 as fundamental rights. Although Great Britain has no formal constitution, the *Magna Carta* of 1215, the *1512 Privilege of Parliament Act*, the *1689 Bill of Rights*, and the *1911 Parliament Act* all protect aspects of the freedom of speech.[2]

Despite such constitutional protections, speech is regulated everywhere. Article 1 of the Canadian Charter of Rights and Freedoms permits "demonstrably reasonable" limitations on freedom of expression. Great Britain and Canada have promulgated legislation regulating racially defamatory speech, group defamation, or speech that incites racial hatred. Although American law is far more protective of speech, some expression is deemed unprotected by the First Amendment, and even protected speech can be regulated in a variety of circumstances.

Accordingly, some governments have enacted laws prohibiting certain content on the Internet and have sought to prosecute users and service providers. Others have tried to control access by insisting on the installation of national *proxy servers* and blocking targeted Websites. Governments also have encouraged *self-regulation* intended to enlist service providers to control customer behavior.

Readers from North America should also understand that in most of the developing world, access to the Internet is primarily through mobile phones. For example, at the time of writing (August 2013) in developing nations in the year 2012, mobile phone subscriptions for Internet access registered at 89.4 percent (5.2 billion subscriptions) of the population, whereas only 6.2 percent of the population (0.4 billion subscriptions) had fixed broadband subscriptions.[3] Thus, restrictions on Internet access include restrictions on mobile-device usage in many countries.

72.1.2 Whose Laws? Whose Standards? It is often thought that governments cannot successfully censor the Internet. After all, one nation's laws may be unenforceable elsewhere. But national laws can be applied to communication intermediaries like Internet service providers (ISPs), blocking access by people in that country.

Moreover, just as speech originating in one nation can reach a worldwide audience, one nation's laws can have powerful effects in other nations. As commentators have put it, "[t]he reach of the Internet multiplies both the number of laws and the number of jurisdictions applicable to speech transmitted on-line."[4] Indeed, "[b]ecause content posted on the Internet is instantaneously transmitted worldwide, Internet users and providers almost automatically face potential liability in any country to which the Internet is connected—not just the nation where the speech was created or where the author lives."[5]

A prominent example is the case in which a French court ordered the United States–based Yahoo! to "take all necessary measures to dissuade and make impossible any access via Yahoo.com to the auction service for Nazi merchandise as well as to any other site or service that may be construed as an apology for Nazism or contesting the reality of Nazi crimes."[6] The French court held that "the simple act of displaying [Nazi artifacts] in France violates Article R645-1 of the Penal Code and therefore [is] a threat to internal public order."[7] It described the mere availability of such information to be "a connecting link with France, which renders our jurisdiction perfectly competent to

rule in this matter."[8] After the French order was handed down, Yahoo! filed suit in the United States seeking a judgment that it need not comply with the French order.[9] The French parties moved to dismiss the U.S. case for lack of personal jurisdiction, but the U.S. district court denied the motion, allowing Yahoo! to proceed.[10] In 2004, however, the U.S. Appellate Court overturned the district court's decision, ruling that the district court judge lacked the authority to hear the case, but also noting that Yahoo! should not expect both to benefit from the worldwide availability of its product and to avoid the possible consequences of that availability.

Resolving international jurisdiction in cyberspace is one of the most difficult and important issues for the Internet, but it is beyond the scope of this chapter. Moreover, the ways in which the laws of one nation may affect Internet speech in another extend beyond the topic of this chapter: state-backed censorship. Even when governments do not themselves target particular kinds of speech, their laws may give private parties legal rights against expression originating overseas, as in the area of defamation.[11] Laws relating to commerce also may raise free speech concerns, as in the areas of intellectual property[12] and privacy regulation. For instance, the 1995 European Union Data Protection Directive generally regulates all *processing* of *personal data*, giving data subjects legal rights quite different from those found in the United States.[13] See Chapter 69 in this *Handbook* for discussion of the Privacy Directive and other regulations affecting the handling of personally identifiable data.

72.1.3 Defining Objectionable Material: International Differences.

For most countries, the major concerns stem from the increased availability of objectionable content over the Internet. This section focuses on how the meaning of *objectionable* differs from country to country.

72.1.3.1 Sexuality.

A common type of objectionable material is speech about sex. In the United States, a series of attempts have been made to regulate the availability of sexually oriented speech over the Internet, although with little success.[14]

Several Middle Eastern nations also block pornographic sites.[15]

In 1995, the Bavarian Justice Ministry informed the value-added network CompuServe that senior company officials could face prison terms for violation of German antipornography laws. German police provided CompuServe a list of the discussion groups that contained potentially objectionable material.[16] In February 1997, Germany indicted Felix Somm, head of CompuServe's German subsidiary, for failure to prevent the dissemination of illegal material on its online service; with this Germany became the first Western democracy to prosecute an official of an online service for illegal Internet content.[17] German prosecutors argued that because CompuServe had access to screening software, the company had the opportunity to block the offending material but failed to do so. Soon thereafter, Germany enacted its *Information and Communications Services Act of 1997* (ICSA). The ICSA, however, has been sharply criticized for failing to resolve the uncertainties faced by ISPs.[18]

72.1.3.2 Hate.

Another major category of objectionable speech is what is loosely called *hate speech*.[19] American hate sites include sites promoting white supremacy, promoting Nazi or neo-Nazi views, supporting skinheads, and promoting the Christian Identity Movement.[20] Internet content that promotes Nazi views and Holocaust denial has attracted considerable international attention.

For instance, the German constitution, known as the *Basic Law*, incorporates freedom of expression as a fundamental individual right. But Article 5, the main provision concerning free speech, expressly provides that the government may limit an individual's expressive right if it conflicts with other people's rights, public order, or criminal laws. Thus, under Germany's general laws and the laws protecting youths, publishing or distributing neo-Nazi or Holocaust denial literature is a criminal offense.[21]

In February 1996, German prosecutors in Mannheim investigated several ISPs, including America Online, regarding alleged distribution of neo-Nazi material on the Internet in violation of anti-Nazi laws.[22] In 2000, German-born Frederick Toben, who used the Website of his Australian-based Adelaide Institute to advocate Holocaust denial, was found guilty of offending the memory of the dead.[23] The lower court had ruled "that German law against inciting racial hatred could not be applied to content on a foreign website."[24] Germany's Federal Court of Justice, the Bundesgerichtshof, overturned the lower court ruling and found that German law applies to Internet content that originates outside the country's borders, so long as people inside of Germany can access that content.[25]

In a 2002 study, Harvard Law School professor Jonathan Zittrain found that Google was systematically excluding search results containing banned content in the French and German versions of its site and that this was apparently "because of pressure applied or perceived by the respective governments."[26]

Canadian law directs courts to balance an individual's free speech rights with societal equality interests. Thus, Canadian courts have upheld hate speech convictions under laws that criminalize the willful promotion of hatred.[27] The Canadian Supreme Court held the statute constitutional.[28] The Court defined hate speech as including a requirement that "the accused 'promote' hatred against an identifiable group, indicating more than simple encouragement or advancement and rather required that the hate monger intend or foresee as substantially certain a direct and active stimulation of hatred against an identifiable group."[29]

The Internet has raised new issues in Canada regarding hate speech. In the 1990s, Bernard Klatt ran the *Fairview Technology Center* out of his home in British Columbia, and about a dozen neo-Nazi, white supremacist, and skinhead clients used his server to publish material against immigration and homosexuality while celebrating Hitler's accomplishments and "Euro-Christianity." The Canadian government attempted to shut down Klatt, but he and his clients insisted that they had not broken any laws.[30] The site suspended its services in 1998.[31] As of the time this chapter was updated (August 2013), the Website consisted of a single page with contact information and had not been updated since 2002.

The Zündelsite, a multilingual archive and news service, may be the most prominent Holocaust-denial Website run by an individual. Its administrator, German Canadian Ernst Zündel, became one of the denial movement's martyrs when he was convicted in Canada for his Holocaust denial pamphlets.[32] The Canadian Supreme Court later overturned his conviction.[33] Zündel's United States–based Web page then became the target of an official inquiry by the Canadian Human Rights Commission.[34] Zündel was deported from the United States to Canada in 2003. The Canadian government determined that he represented a risk to national security, and deported him to Germany, where he was sentenced in 2007 for the maximum term of five years on outstanding charges of incitement for Holocaust denial but was released in March 2010.[35]

72.1.3.3 Politics and Culture. Various countries also regulate Internet content that generally threatens their government or political, social, and cultural values.

72.1.3.3.1 People's Republic of China. At the end of 2006, there were approximately 137 million Internet users in China, an increase of 34 million over the previous year.[36] The Chinese government, although eager to capitalize on the Internet, feels that the "Internet is a very real threat to the hold their dictatorial regime has over the country."[37] Accordingly, China severely restricts communication via the Internet, including much dissent and the free reporting of news. The *Measures for Managing Internet Information Services* prohibit private Websites from publishing news without prior approval from Communist Party officials.[38]

A major obstacle for would-be Chinese Internet users is the registration process required to get online. One must pick an ISP and provide identification. One must also fill out a *Police File Report Form* that will be sent to the ISP, the local *Public Security Bureau* (PSB), and the provincial-level PSB *Computer Security and Supervision Office*. In addition, users must complete a Net Access Responsibility Agreement, pledging "not to use the Internet to threaten state security or reveal state secrets" or to read, reproduce, or transmit material that "endangers the state, obstructs public safety, or is obscene or pornographic."[39] The ISP application itself requests information such as employer's address, home address, and profession; home, office, cell phone, and pager numbers; as well as details about the computer, the modem type, and permit number, which is assigned by the PSB.[40]

The following news summaries from 2001 were posted on the *Digital Freedom Network* site, which no longer has any of the original reports online; however, the original article is available on the Internet Archive.[41]

- In January 2000, the Chinese newspaper the *People's Daily* published new Internet regulations issued by China's State Secrecy Bureau. These laws "ban the release, discussion, or transfer" of "state secret information" on bulletin board systems, chat rooms, or Internet news groups. Email users were banned from sending or forwarding state secrets by email. In addition, all Websites are required to undergo a security check. A provision requiring that anyone using encryption technology register with the government was later rescinded.
- In May 2000, after the *China Finance Information Network*, a financial Website, republished a Hong Kong newspaper about corruption by a provincial official, Chinese public security officials suspended the site for 15 days and fined the site 15,000 yuan (US$1,800). On August 3, 2000, state security officials forced a Chinese ISP to shut down the Xinwenming (New Culture Forum) site run by a group of dissidents in Shandong Province for posting "reactionary content."
- The following month, China's State Council passed Internet laws that limited foreign investment and imposed strict surveillance requirements on Internet content by mandating that Internet content providers monitor both the material they publish and the people who access that material, including the times users log onto the Internet, their account numbers, their Internet addresses, and the phone numbers they are dialing in from.
- On November 7, 2000, the Ministry of Information Industry and the Information Office of the State Council released new laws banning commercial media organizations from setting up independent news sites; prohibiting commercial portals from carrying news items based on their own sources; and requiring China-based Websites to obtain permission from the State Council Information Office before linking to overseas news sites or carrying news from overseas news media. The new Internet laws also restrict the content of online chat rooms and bulletin boards.

- In December 2000, the National People's Congress passed a new law to foster Internet safety. This law criminalizes several forms of online political activity, including using the Internet to "incite the overthrow of state power, topple the socialist system, or … destroy national unity," promote "cults," or support the independence of Taiwan.
- The Ministry of Public Security announced on February 27, 2001, that it released Internet filtering software called *Internet Police 110* on February 26. The software comes in three versions for households, Internet cafés, and schools that can monitor Web traffic and delete or block messages from sources deemed "offensive." According to a news report published in the China Securities News, it can send text or voice warnings to network administrators about any unauthorized Internet surfing.
- On March 20, 2001, official Chinese media reported that China is developing a surveillance system to monitor activities on the Internet, which will be similar to the data-recording "black box" installed in commercial airplanes. According to Hong Kong *iMail*, it will be able to record all communications through the Internet.
- In April 2001, the official Chinese news agency Xinhua reported that China would impose a three-month ban on the opening of new Internet cafés. This is part of "a major offensive against unchecked use of the Internet. Xinhua said authorities are to conduct a massive probe into existing Internet outlets, which it views as potential hotbeds of dissent and vice."
- In addition, after January 2000, the Chinese government began arresting individuals who use the Internet in a manner they deem dangerous, often for online political or religious activity. Qi Yanchen, sentenced to four years in prison on September 19, 2000, is the first Chinese convicted of subversion for material he wrote that was published on the Internet. Qi was officially charged for writing articles in the May 6, 1999, and May 17, 1999, United States–based Chinese dissident email publication *Dacankao* (V.I.P. Reference). Guo Qinghai, a freelance writer, was arrested in September 2000 for "subverting state power." Guo published articles on the Internet that discussed Qi Yanchen's case. He also posted, on overseas online bulletin boards, essays promoting political reforms in China. He was sentenced to four years in prison.
- Zhang Haitao, creator of the only China-based Website on the outlawed Falun Gong, was charged with subversion on October 11, 2000, and is accused of establishing a site promoting Falun Gong and of posting an online petition urging followers to protest the government ban on the group. Zhang Ji, a college student in Heilongjiang Province, was charged on November 8, 2000, with "disseminating reactionary documents via the Internet." Authorities say Zhang had emailed information to United States– and Canada-based Websites of the Falun Gong religious group as well as downloaded news about the group and shared it with others in China.

Google has had a rocky relationship with the Chinese government, with its entire service being blocked several times for serving information critical of government policy. In January 2006, Google agreed to block a list of search terms that would be provided by the Chinese government and limit searches from within China to this partial set of results at google.cn. The independence of Taiwan and the 1989 Tiananmen

Square massacre remain on the list of forbidden topics. Given the history of outright bans, Google seems to have decided that some content was better than none.[42]

In May 2011, a new body, the State Internet Information Office (SIIO), was created to coordinate state control over all aspects of Internet usage. Within weeks, the SIIO had shut down over 6,000 sites offering reputation-management services (some of them allegedly fraudulent).[43]

In December 2012, reports from Beijing indicated that "China unveiled tighter Internet controls on Friday, legalizing the deletion of posts or pages which are deemed to contain 'illegal' information and requiring service providers to hand over such information to the authorities for punishment." A Reuters report included this summary:

> "Service providers are required to instantly stop the transmission of illegal information once it is spotted and take relevant measures, including removing the information and saving records, before reporting to supervisory authorities," the rules state.
>
> The restrictions follow a series of corruption scandals amongst lower-level officials exposed by Internet users, something the government has said it is trying to encourage.
>
> Li Fei, deputy head of parliament's legislative affairs committee, said the new rules did not mean people needed to worry about being unable to report corruption online. But he added a warning too.
>
> "When people exercise their rights, including the right to use the Internet, they must do so in accordance with the law and constitution, and not harm the legal rights of the state, society … or other citizens," he told a news conference.
>
> Chinese Internet users already cope with extensive censorship measures, especially over politically sensitive topics like human rights and elite politics, and popular foreign sites Facebook, Twitter, and Google-owned YouTube are blocked.
>
> Earlier [in 2012], the government began forcing users of Sina's wildly successful Weibo microblogging platform to register their real names.[44]

The *New York Times* maintains an extensive and up-to-date compilation of its articles about Internet censorship in China.[45]

72.1.3.3.2 Islamic Countries. According to *Human Rights Watch*, the United Arab Emirates (UAE) is "the regional leader in advocating censorship of the Web through the use of high-tech means."[46] Dial-up users in the UAE do not access the Internet directly; they dial into a proxy server maintained by *Etisalat*, the state telecommunications company, in collaboration with a U.S. firm that is contracted to maintain and update the software. Although UAE officials have insisted that the proxy server's sole purpose is to block pornographic sites, at least one nonpornographic site, that of the *Gay and Lesbian Arabic Society* (www.glas.org) is blocked. When asked about the site, an Information and Culture Ministry official explained that there had been complaints about it.

Saudi Arabia bans publishing or even accessing various types of online expression, including "[a]nything contrary to the state or its system"; "[n]ews damaging to the Saudi Arabian armed forces"; "[a]nything damaging to the dignity of heads of states"; "[a]ny false information ascribed to state officials"; "[s]ubversive ideas"; and "[s]landerous or libelous material."[47]

All 30 of the country's ISPs are linked to a ground-floor room at the Riyadh Internet entranceway, where all of the country's Web activity is stored in massive cache files and screened for offensive or sacrilegious material before it is released to individual users.

The central servers are configured to block access to "sensitive" sites that might violate "the social, cultural, political, media, economic, and religious values of the Kingdom."[48]

According to Human Rights Watch, although official Saudi explanations of Internet censorship focus on materials deemed offensive to conservative Muslim sensibilities, Saudi blocking apparently extends to political sites. In early 1999, the site of at least one exiled dissident group, the Committee against Corruption in Saudi Arabia (www.saudhouse.com), was reportedly blocked.

Yemen's telecommunications service provider, *Teleyemen*, told Human Rights Watch that there is a "general requirement … to limit information which is considered to be undesirable in terms of causing offenses against social, religious, or cultural standards" (Eric Goldstein, The Internet in the Mideast and North Africa: Free Expression and Censorship, p. 24) and that Teleyemen uses the Surfwatch censorware program in conjunction with a proxy server.

More recently, controversies have erupted when Middle Eastern and other Islamic groups or nations have protested content offensive to Islam that was hosted outside their own borders. In 2005, a Danish newspaper published caricatures of the prophet Muhammad, which many Muslims found to be offensive. The cartoonist, the newspaper, and the Danish people as a whole were targeted with boycotts and death threats. Imams in Denmark brought suit to have the government punish the newspaper but had no success in the courts. In 2007, Turkey briefly blocked all access to YouTube in protest of a video insulting the nation's founder, Mustafa Kemal Atatürk. Google, which had purchased YouTube earlier that year, quickly removed the video in question, and Turkey revoked the ban.[49]

A report from *Freedom House* in 2012 analyzing Internet censorship around the world listed 47 countries; notably, predominantly Muslim countries differed in the direction of their controls, with some becoming more restrictive and other becoming more liberal. Highlights from the organization's press release[50] and the full report[51] (662 pages) include detailed, country-by-country reports spanning several pages each and including extensive reference material. Additional commentary on countries of interest in this section includes:

- **Bahrain**: "Following the February 14, 2011, protests, the government intensified censorship and surveillance of advanced Web 2.0 applications and blocked interactive exchanges online, particularly when its political agenda was not supported. Internet connections became very slow, making it difficult to upload media, and some locations were entirely offline. Since February 2011, most live broadcasting websites that were popular among protesters have been blocked. PalTalk, a chatting service that was used to conduct political seminars with prominent guests and mass online audiences, has been blocked since June 2011, while many blogs critical of government views were also blocked in 2011, particularly those that documented the protests and government crackdown. …"
- **Egypt**: "The Supreme Council of the Armed Forces (SCAF) maintained many of its predecessor's tactics of internet control, while intensifying others. Mobile phones, the internet, and social media remained under vigorous surveillance, bandwidth speeds were throttled during specific events, and SCAF-affiliated commentators manipulated online discussions. Several activists and bloggers were intimidated, beaten, shot at, or tried in military courts for "insulting the military power" or "disturbing social peace." Despite recent elections, the future trajectory of internet freedom in Egypt remains precarious and uncertain."

- **Ethiopia**: "In 2011, in the wake of the Arab Spring protests in the Middle East and several online calls for similar demonstrations in Ethiopia, … the government reacted by strengthening internet censorship and carrying out a systematic crackdown on independent journalists, including at least one blogger. Beginning in June 2011, over ten journalists were sentenced to long prison terms, … mostly on questionable charges of terrorism. Among them was the editor of an exiled online news Website who was sentenced in absentia to life imprisonment. A prominent dissident blogger based in Ethiopia was also arrested in September 2011 and sentenced to 18 years in prison in July 2012. … The latest crackdown is part of a broader trend of growing repression against independent media since the 2005 parliamentary elections, in which opposition parties mustered a relatively strong showing. …"
- **Indonesia**: "Since 2010, internet and mobile phone usage has continued to grow at a fast pace. Meanwhile, the popularity of social-networking applications has increased exponentially, with Indonesia becoming home to some of the largest contingents of Twitter and Facebook users in the world. The authorities have subsequently sought to regulate online content, citing fears of the internet's use for the spread of pornographic, blasphemous, and terrorism-related content. In the process, a number of actions taken, including passage of the Law on Information and Electronic Transactions (ITE Law) of 2008, have fallen short of international democratic standards. This trend continued in 2011 with the passage of the State Surveillance Law and increased filtering of content loosely defined as sexually explicit. Meanwhile, a surprising criminal conviction under the ITE Law of a housewife who disseminated critical information about a hospital and a series of other questionable criminal cases filed under the ITE exacerbated the atmosphere of legal uncertainty surrounding freedom of expression online and raised concerns of greater restrictions on internet freedom in the future."
- **Iran**: "The Iranian authorities used more nuanced tactics in a continued campaign against internet freedom that began after disputed elections in 2009. These tactics included: upgrading content filtering technology, hacking digital certificates to undermine user privacy, and moving closer to establishing a National Internet. Iranian judicial authorities also meted out some of the harshest sentences in the world for online activities, including imposing the death penalty on three bloggers and IT professionals."
- **Jordan**: "Jordan, a small kingdom of about six million people, prides itself on offering relatively broad freedom to use the internet and officially blocks only one Website. The Jordanian government's response to public protests in 2011 was relatively mild compared to neighboring countries. Moreover, the king and parliament passed a set of constitutional amendments that could improve human rights protections and free expression. Nevertheless, restrictions on internet freedom exist and have increased since January 2011. News Websites, which have become a vital source of information in a country where traditional media freedom is limited, often face pressure from state actors to delete articles deemed politically sensitive. In April 2012, a government council passed a resolution that could require online news Websites to register with the government, a policy that if implemented risks curbing their independent reporting and analysis. Meanwhile, surveillance, physical attacks against bloggers and online journalists, and hacking attacks against prominent news Websites also present a threat."

- **Libya**: "The political unrest and armed conflict that occurred in Libya in 2011, which resulted in the death of Muammar al-Qadhafi after over 40 years in power, led to a dramatic regime change. The country shifted from being one of the world's harshest dictatorships to a postconflict aspiring democracy. As of May 2012, the government of Libya was comprised of a National Transitional Council (NTC) formed during the conflict and an appointed interim government mandated to steer the country towards elections scheduled for July 2012, after which a new constitution will be drafted. These political changes were also reflected in the internet freedom landscape. As such, this report straddles three radically different periods: a highly restrictive environment under Qadhafi, a partial internet and telephone blackout for much of 2011, and a relatively open online information landscape since the rebel victory in October 2011."
- **Pakistan**: "Disconcerting recent developments in Pakistan include a ban on encryption and virtual private networks (VPNs), a death sentence imposed for transmitting allegedly blasphemous content via text message, and a one-day block on all mobile phone networks in Balochistan province. Several other initiatives to increase censorship—including a plan to filter text messages by keyword and a proposal to develop a nationwide internet firewall—were officially shelved in response to civil society advocacy campaigns, although some suspect that the government is still working on them behind closed doors."
- **Saudi Arabia**: "While the recent wave of protests and demonstrations in the Kingdom of Saudi Arabia has not changed the political landscape in the country as in other Arab countries such as Tunisia, Egypt, and Libya, the internet freedom landscape in the Kingdom has no doubt changed considerably over the past year. Inspired by the Arab Spring events in 2011, millions of people in Saudi Arabia flooded social media sites such as Facebook, YouTube, and Twitter, and hundreds, if not thousands, participated in political campaigns to demand political, social, and economic reforms, fostering the emergence of innumerable political activists on social media. However, as soon as Saudis went online to express their opinions, demand actions, and organize demonstrations, the Saudi government took immediate steps to respond to what it regarded as a national security threat. The government issued warnings banning protests, reminded people via email of the ban, and threatened the youth through the BlackBerry multimedia message service (MMS) to discourage them from participating in demonstrations. The authorities also detained and intimidated hundreds of online political activists and online commentators, implemented strict filtering mechanisms to block sensitive political content from entering the Saudi internet, recruited thousands of online supporters to warn against the call for protests and demonstrations as a counter measure, and continued to apply its excessive monitoring of internet users."
- **Syria**: "The Syrian government engages in extensive filtering of Websites related to politics, minorities, human rights, and foreign affairs, and such censorship has expanded in recent years. Tests conducted by the OpenNet Initiative (ONI) in 2008–2009 found pervasive blocking of Websites related to opposition to the Assad regime, human rights groups, the Muslim Brotherhood, and activism on behalf of the Kurdish minority. A range of Websites related to regional politics were also found to be blocked, including several Lebanese newspapers and the Websites of groups campaigning to end Syrian influence in Lebanon, as well as the prominent London-based Arabic newspapers Al-Quds al-Arabi and Al-Sharq

> al-Awesat. Access to the entire Israeli domain name ".il" was also restricted. Internet censorship continued to worsen in 2011 and early 2012. Activists and average users consistently complained about the extensive and unprecedented blocking of circumvention tools, internet security software, and applications that enable anonymous communications. Websites used to mobilize people for protests or resistance against the regime, including those of the network of Local Coordination Committees (LCCs) that emerged as the uprising progressed, were blocked. Online initiatives to gather information and raise public awareness, such as the Mondaseh Website, were blocked as well. However, the Websites of most international news sources and human rights groups have remained accessible."

72.1.3.3.3 Burma (Myanmar). In 1988, a new military clique called the *State Law* and *Order Restoration Council* (SLORC) took power in Burma, now known as *Myanmar*. They managed to achieve cease-fires with 15 armed organizations, but they refused to step down, asserting that only "the military can ensure national unity and solidarity,"[52] even after the National League for Democracy's overwhelming victory in the 1990 election. In 1997, the SLORC was renamed the State Peace and Development Council (SPDC), but the membership remained largely unchanged.

SLORC/SPDC maintains power partly by controlling access to information. In January 2000, the Myanmar Post and Telecommunication (MPT), then the only ISP in Myanmar, passed a measure that forbids Internet users from posting political speech on the Web, such as writings "detrimental to the interests of the Union of Myanmar" or "directly or indirectly detrimental to the current policies and secret security affairs of the government of the Union of Burma."[53] Only the person granted an Internet account may use it, and the account holder is held responsible for all use of that account. Web pages may not be created without government permission, and violation of any regulation is subject to legal action.[54] In 2005, Myanmar's only private ISP, Bagan Cybertech, was nationalized under the control of the MPT, and the government has entered into an agreement with U.S. Web content filtering company Fortinet in order to better control Web access from within Myanmar.[55]

The Burmese military has been effective at "warding off any Internet revolution by limiting the rights of its citizens."[56] The Freedom House report cited in Section 72.1.3.3.2 identifies persistently weak infrastructure for both standard computer-based Internet connectivity and for mobile-phone access to the Internet as barriers to Internet access and usage.[57]

72.1.3.3.4 Russia. The Freedom House report[58] describes the situation in Russia at the end of 2012 as follows:

> After independent television channels were eliminated and press regulations tightened from 2000–2001, the internet became Russia's last relatively uncensored platform for public debate and the expression of political opinions. In response, the government has tried various tactics to suppress citizens' right to free speech online over the years. In 2009–2010, many bloggers were harassed and opposition blogs were hacked. While these incidents still occurred in 2011, the tactics to restrict freedom of expression online have slightly changed, with the deployment of distributed denial-of-service (DDoS) attacks … and various smear campaigns … to discredit online activists becoming more common. Extralegal intimidation of social network activists and independent forum moderators … has become another line of pressure over the online world through strategies such as informal meetings with the security services, calls from the Federal Security Service (FSB) to the parents of activists, … or the sudden refusal of forum ad sponsors to buy advertisements.

In May 2013, a wide range of Websites were blocked in central Russia after a district court responded to a prosecutor's claim that they were making available instructions on how to bribe officials. However, none of the Website owners were informed of the decision, some of the sites were for independent news organizations (rather than state-controlled outlets), and the blockage even included information such as employment advertisements. Activists described the incident as "Internet censorship on a whim."[59]

72.2 U.S. CONTEXT: FIRST AMENDMENT RIGHTS. The First Amendment of the U.S. Constitution strongly protects freedom of speech in the United States. In several ways, the right to speak is the most powerful of all U.S. constitutional rights. First Amendment scrutiny can be much stricter than in other areas of constitutional law. It is usually easier to get in front of a court when "speech" is at issue, and to win, because many technical legal requirements are relaxed.

Thus, although the right to speak is not constitutionally absolute, and may be restricted by government, several basic principles limit governmental authority to restrict freedom of speech, more than it is limited from restricting other activities. These principles have been articulated through a variety of legal doctrines.

72.2.1 What Does the First Amendment Protect? Today, the First Amendment generally protects all forms of communicative or expressive activity[60] from "state action" or governmental interference.[61] Courts generally presume that linguistic acts or visual depictions are covered by the First Amendment but require litigants to show that nonlinguistic acts are speech unless the law has already so held.[62]

Nevertheless, the coverage of the First Amendment is quite broad. It includes not only the right to speak but also the right *not* to speak, or against "compelled" speech,[63] the right to speak or associate anonymously,[64] and the right to read or to receive information.[65] Often, activity that is not necessarily communicative is protected because it enables or facilitates expression. Highly discretionary city licensing of newsstands was found to be an invalid scheme of licensing that affected liberty of circulation.[66] Money is not speech, but election campaign contributions and expenditures are protected as speech.[67]

In addition, the First Amendment requires that the government observe some degree of neutrality in its treatment of speakers[68] and speech media[69]; it cannot arbitrarily prefer one speaker or one medium to another.

72.2.2 Basic First Amendment Principles. There is no consensus or accepted theory underlying U.S. protection of free speech. The different rationales for freedom of speech can be categorized as more or less "positive" or "negative." Positive justifications identify speech as related to some special moral, social, or political value. For example, the "marketplace of ideas" theory is based on the notion that "the best test of truth is the power of the thought to get itself accepted in the competition of the market."[70] Courts also have defended the right to free expression as embodying the value of "speaker autonomy,"[71] thus "putting the decision as to what views shall be voiced largely into the hands of each of us, in the hope that use of such freedom will ultimately produce a more capable citizenry and more perfect polity, *and in the belief that no other approach would comport with the premise of individual dignity and choice upon which our political system rests*."[72]

Negative theories presume the value of speech and focus instead on the dangers of government regulation of speech. The main idea here is that speech is as vulnerable to threats of punishment as to punishment itself. "It is characteristic of freedoms of

expression in general that they are vulnerable to gravely damaging yet barely visible encroachments."[73] This fear of "chilling effects" on speech sometimes is associated with the possibility that juries may not treat unpopular speakers or unorthodox ideas fairly, but it also reflects a concern that individuals may censor themselves. Such self-censorship can harm public debate while being almost "invisible" to the courts because it occurs through many small individual choices.[74]

72.2.2.1 Harm/Causation. Perhaps the most important First Amendment principle is that the government has the burden to show that the speech it wishes to regulate truly causes significant harm. This principle was not accepted for many years; in the early days of First Amendment law, speech could be punished as an "attempt" if the natural and reasonable tendency of what was said would be to bring about a forbidden effect. Thus, for instance, individuals who had mailed circulars to draftees arguing that conscription was unconstitutional and urging them to assert their rights were convicted for conspiring to attempt to obstruct the draft.[75]

Today, courts are far more skeptical about government claims of harm. If speech could be restricted merely because the mathematical expectation of harm were significant, speech would be easier to restrict as it were more widely disseminated. Such a result is inconsistent with the First Amendment's protection of open, public discourse. As a classic opinion stated, "no danger flowing from speech can be deemed clear and present, unless the incidence of the evil apprehended is so imminent that it may befall before there is opportunity for full discussion. If there be time to expose the evil by the processes of education, the remedy to be applied is more speech, not enforced silence."[76]

The concern here is not only for the speech itself; it includes a concern that the speaker can properly be held responsible for harm. The modern approach is exemplified by a case where a speaker was convicted for advocating violence. The Supreme Court reversed the conviction because the law did not require the government to show both that the speaker intended to produce "imminent lawless action" and that the speech was "likely to incite or produce such action."[77] The requirement of intent helps protect speakers from being held responsible for the acts of others. As one court put it, "Much speech is dangerous. Chemists whose work might help someone build a bomb, political theorists whose papers might start political movements that lead to riots, speakers whose ideas attract violent protesters, all these and more leave loss in their wake."[78] Without careful attention to harm and causation, the right to speak would mean little.[79]

72.2.2.2 Neutrality. Not all reasons that government might use to restrict speech are valid. Thus, a second basic First Amendment principle is governmental neutrality as to speech content. Government is likely to have illegitimate reasons for wishing to restrict speech, such as the desire to suppress political criticism, but it may be able to disguise such hostility by claiming that less selfish interests, such as "public order" or the civility of discourse, are truly at stake.[80] Also, because public officials may have a strong self-interest in stifling dissent, they are likely to ignore or undervalue the social interest in free speech. Such motivations may lead to regulation that skews or distorts public debate. In general, then, government may not regulate speech based on its substantive content or the message it conveys.[81]

72.2.2.3 Precision. Third, all speech regulation must be precise. "Because First Amendment freedoms need breathing space to survive, government may regulate in the area only with narrow specificity."[82] Much of the First Amendment's force stems

from this principle. Imprecise regulation not only chills speech but also gives officials too much discretion to discriminate on the basis of content. Without clearly stated standards, courts cannot easily ferret out content-based discrimination.[83] Insisting on precision also allows courts to avoid assessing the importance of the legislature's goals; they need only say that the means are too blunt.

72.2.2.4 Background Doubt-Resolution Principles. Adding to the strength of First Amendment protection for speech are several principles for resolving uncertainty. Speech regulation generally requires knowledge or intent on the speaker's part, which protects bookstores and others that "carry" others' speech.[84] Moreover, the factual burden of proof is usually on the person or entity seeking to restrict speech.[85] Finally, speech that in itself might be disfavored receives "strategic" protection in order to protect undue chill to public discourse; for instance, while false statements of fact may be considered of little First Amendment value, permitting the punishment of mere falsehood would overly burden speakers.[86] Put simply, doubts generally are called in favor of speech.

72.2.3 Limitations on Government Interference with Speech

72.2.3.1 Distinction between Substantive Tests and Procedural Safeguards: Example of Prior Restraints. A critical distinction in First Amendment (and other) law is between substantive and procedural review. The difference is best explained by reference to "prior restraints," which are presumed unconstitutional.

Most people are familiar with prior restraints as judicial orders preventing speech or publication. In such cases, government "carries a heavy burden of showing justification for the imposition of such a restraint."[87] For instance, the Supreme Court required the government to demonstrate that publication of the *Pentagon Papers* "will surely result in direct, immediate, and irreparable damage to our Nation or its people."[88] Judicial "gag orders" on press publication are subject to similarly searching scrutiny.[89] In these cases, courts are concerned with whether the particular restraint on speech is sufficiently justified.

A second line of prior restraint doctrine illustrates procedural scrutiny. American hostility to prior restraints is a reaction to English press licensing systems under which nothing could be printed without the prior approval of government or church authority. The fear is that discretionary permit schemes make freedom of speech "contingent upon the uncontrolled will of an official."[90] Thus, the law authorizing such a licensing scheme must provide "narrowly drawn, reasonable and definite standards for the [administering] officials to follow."[91] Thus, licensing schemes, such as for screening obscene movies, are constitutional only if accompanied by procedural safeguards necessary "to obviate the dangers of a censorship system."[92] There must be a "prompt final judicial decision" reviewing any "interim and possibly erroneous denial of a license." Moreover, "because only a judicial determination in an adversary proceeding ensures the necessary sensitivity to freedom of expression," the censor must, "within a specified brief period, either issue a license or go to court."[93] In such cases, the particular restraint is relatively unimportant: The issue is whether the statute itself contains sufficient procedural safeguards.

72.2.3.2 Levels of Substantive Scrutiny and the Issue of Content Neutrality. First Amendment cases typically feature three levels of substantive scrutiny, turning mainly on how the statute regulates the meaning of speech: by subject, by

viewpoint, or without reference to meaning at all. In general, strict scrutiny applies to laws that attempt to stifle "speech on account of its message."[94] In such situations, the government must demonstrate both that there is a compelling interest in restricting the speech and that it has chosen the least restrictive means of furthering that compelling interest.

When the government targets not subject matter but particular views taken by speakers on a subject, the violation of the First Amendment is all the more blatant.[95] Viewpoint discrimination is thus an egregious form of content discrimination. The government must abstain from regulating speech when the specific motivating ideology or the opinion or perspective of the speaker is the rationale for the restriction.[96] Such scrutiny is so stringent that it applies to otherwise "unprotected" speech.[97]

However, the government often regulates speech without reference to its meaning or communicative impact. So-called content-neutral speech regulation is subject to intermediate rather than strict scrutiny "because in most cases they pose a less substantial risk of excising certain ideas or viewpoints from the public dialogue."[98] Such laws "do not pose such inherent dangers to free expression, or present such potential for censorship or manipulation, as to justify application of the most exacting level of First Amendment scrutiny."[99]

Intermediate scrutiny comes in two main flavors. Content-neutral regulation of the time, place, or manner of speech must be narrowly tailored to serve a significant government interest and leave open ample alternative channels for communication.[100] Content-neutral regulation in other contexts must further a substantial government interest and not burden substantially more speech than necessary to further that interest.[101]

A third type of intermediate scrutiny is the "secondary effects" doctrine, which operates as an exception to the general rule that content-based speech regulation is subject to strict scrutiny. The doctrine asks whether the harm is attributable to the communicative aspects of the speech; if not, the regulation is said to be aimed not at speech but rather at its "secondary effects," and the regulation is deemed content-neutral.[102] Thus, city zoning ordinances have been upheld even though they applied only to theaters that showed adult films because their purposes of preventing crime and protecting property values and the quality of urban life were deemed unrelated to speech content. This doctrine has been criticized for permitting speech suppression "whenever censors can concoct 'secondary' rationalizations for regulating the content of political speech."[103] The main limit on the doctrine is the rule that "listeners' reaction to speech is not a content-neutral basis for regulation."[104]

72.2.3.3 Types of Procedural Scrutiny: Vagueness and Overbreadth.

Laws generally may not "set a net large enough to catch all possible offenders, and leave it to the courts to step inside and say who could be rightfully detained, and who should be set at large."[105] This hostility to vague laws is greater when speech is concerned; courts will assume that ambiguous legal language will be used against speech.[106] The void-for-vagueness doctrine emphasizes both fair notice to citizens of what conduct is prohibited and clear standards for law enforcement to follow.[107]

A closely related doctrine is that of overbreadth. A law is overbroad when it "does not aim specifically at evils within the allowable area of [government] control, but … sweeps within its ambit other activities that constitute an exercise" of First Amendment rights.[108] The danger is "not merely the sporadic abuse of power by the censor" but the "continuous and pervasive restraint on all freedom of discussion that might be reasonably regarded as within its purview."[109] Accordingly, the overbreadth doctrine permits facial[110] invalidation of laws that inhibit speech if impermissible

applications of the law are substantial, when "judged in relation to the statute's plainly legitimate sweep."[111] Both doctrines exemplify the basic precision principle.

72.2.4 Exceptions Where Speech Can Legally Be Limited. The government can restrict speech both directly and indirectly. Roughly speaking, the government is more limited in direct regulation of speech than in indirect regulation. First Amendment law also contains many subdomains with their own special rules, such as over-the-air broadcasting,[112] government enterprises,[113] and government subsidies.[114]

72.2.4.1 Neutral Laws of General Applicability. Many laws, in some sense, restrict speech to an extent. Absent special circumstances, "generally applicable" laws can be applied to speech without triggering First Amendment scrutiny. Classic examples are laws against speeding applied to a news reporter hurrying to file a story or taxes that apply to books and newspapers as well as other goods.

The "generally applicable" doctrine contains its own limits. Courts will look to see whether the law appears to be a pretext for targeting speech and whether it was based on perceived secondary effects.[115] If so, then there will be some level of First Amendment scrutiny.

72.2.4.2 Constitutionalization of Crimes and Torts: Advocacy of Illegality, Threats, and Defamation. Another major category of First Amendment doctrines relates to crimes and torts directed at speech. In general, the First Amendment does not prevent the application of criminal law to speech simply because it is speech. Although agreeing to participate or assist in doing an unlawful act is usually accomplished through "speech," the larger course of conduct in which that speech is integrated overrides any potential First Amendment applicability.[116] Even so, there can be concern for speech-suppressive effects.

Criminal laws that are likely to involve speech "must be interpreted with the commands of the First Amendment clearly in mind."[117] As discussed in Section 72.2.2.1, laws that criminalize incitement to unlawful action must require proof of both intent and harm. Similarly, threatening may be criminally punished if the law is carefully defined and applied to "true threats" but not "political hyperbole."[118] Recently, the Supreme Court held that federal law prohibiting the disclosure of unlawfully intercepted communications could not be applied to those who had innocently acquired the information of public concern.[119]

Tort laws that target speech also must be construed and applied in light of First Amendment concerns. The prime example is libel or defamation law. For many years, the First Amendment simply did not apply to libel suits. In 1967, however, the Supreme Court held that libel suits were a form of government action because courts must enforce libel judgments, and thus imposed First Amendment restrictions on such lawsuits.[120]

72.2.4.3 Unprotected Speech. Certain categories of speech, such as *obscenity*, *child pornography*, *fighting words*, *sedition*, *defamation*, and so on, are considered unprotected even though they are linguistic or pictorial. Here the issue is not that the speech is not communicative but that its value is low; *fighting words*, for instance, are deemed "utterances [that] are no essential part of any exposition of ideas, and are of such slight social value as a step to truth that any benefit that may be derived from them is clearly outweighed by the social interest in order and morality."[121]

Even so, the definition of *fighting words* is closely limited. First, the words must "by their very utterance inflict injury or tend to incite an immediate breach of the peace."[122] Second, they must be "directed to the person of the hearer."[123]

The Supreme Court defines obscenity as works that, (1) taken as a whole, appeal to the prurient interest of the average person in sex, (2) portray sexual conduct in a patently offensive way, and (3) taken as a whole, lack serious literary, artistic, political, or scientific value.[124] This definition illustrates how the Supreme Court has sought to ensure that even laws that target unprotected speech do not reach too far. In particular, the first prong of the obscenity test is evaluated with reference to "contemporary community standards." As the Supreme Court observed, "It is neither realistic nor constitutionally sound to read the First Amendment as requiring that the people of Maine or Mississippi accept public depiction of conduct found tolerable in Las Vegas, or New York City."[125] Moreover, any statute regulating obscenity must define the "sexual conduct" at issue.

Child pornography is defined generally as works that visually depict sexually explicit conduct by children below a specified age.[126] As with obscenity, the statute must clearly define the conduct that may not be depicted[127] and must require some element of knowledge or intent. Unlike obscenity, however, legally regulated child pornography need not appeal to the prurient interest of the average person, need not portray sexual conduct in a patently offensive manner, and need not be considered as a whole.[128] Child pornography is the least protected type of speech. Note that *virtual* child pornography (i.e., child pornography with adults representing children or with digital modifications to give the impression that children are involved) is not regulated. The U.S. Supreme Court clarified the distinction, noting that in actual child pornography, the perpetrator commits child abuse in the creation of the pornography; no such abuse is present in the virtual version.[129]

But even unprotected speech is not entirely invisible to the First Amendment. Regulations that target such speech must be based on the reason why the speech is categorized as unprotected; thus, even though fighting words may be regulated because of their potential to generate violence, regulating only "fighting words" that relate to race, color, creed, religion, or gender would "impose special prohibitions on those speakers who express views on disfavored subjects," amounting to unconstitutional viewpoint discrimination.[130]

72.2.4.4 Less-Protected Speech, Media, and Recipients. Other categories of speech are protected by the First Amendment but may be regulated more easily. The main examples are commercial speech, broadcast speech, and speech that is *indecent* or *harmful to minors*.[131]

Commercial speech is defined as speech that proposes a commercial transaction; the main example is advertising. Originally, such speech was considered unprotected, but today its regulation is subject to a form of intermediate scrutiny.[132]

Broadcast speech is also subject to a form of intermediate scrutiny.[133] The Supreme Court has long adhered to the view that different media of expression may require different First Amendment analyses. The standard for regulating the speech of cablecasters is not entirely clear.[134]

With regard to sexual content, there are three forms of less-protected speech. First, although nonobscene sexual expression is in a sense "fully" protected, cities may rely on "secondary effects" analysis to regulate the location of bookstores and theaters that sell "adult" material or offer "adult" entertainment like nude dancing.[135]

Second, the Supreme Court has permitted substantial regulation of indecent speech over the airwaves[136] and on cable TV.[137] Indecency regulation is closely tied to the interest of preventing minors' access to sexual material. As that interest varies, so does the latitude for regulation. For example, the regulation of indecent speech delivered telephonically was not permitted.[138] Thus, indecency is an odd speech category that is regulated primarily with respect to the audience: The government has a strong interest in regulating minors' access to indecency, but it has no such interest in regulating adults' access.

Closely related to the category of *indecency* is materials judged *harmful to minors* (HTM). Roughly speaking, this category is like obscenity, but geared to those under the age of 17.[139] This category has become increasingly important in the censorware context given the full protection accorded to speech over the Internet. Courts have narrowly tailored HTM-based access restrictions to their constitutional boundaries. This point was made clear in *Reno v. ACLU*, where the Supreme Court held unconstitutional a general prohibition of indecent speech on the Internet. The Court found:

- Parents may disseminate HTM speech to their children.[140]
- The HTM concept applies only to commercial transactions.[141]
- The government may not simply ban minors' exposure to a full category of speech, such as nudity, when only a subset of that category can plausibly be deemed HTM.[142]
- The government interest is not equally strong throughout the HTM age range.[143]

72.2.5 Legislation and Legislative Initiatives in the United States. Since the advent of the Internet, there have been many attempts to regulate access to Internet speech. Such legislation can be divided into two basic categories. In the first category, the law imposes liability for communicating certain kinds of material over the Internet. Such laws did not require the use of software to block Internet sites or to filter their content because they focused on the conduct of the speaker or publisher. In the second category, the law requires or encourages the use of filtering software by those who make Internet content available.

In 1996, Congress enacted the *Communications Decency Act* (CDA), which sought to regulate the communication of "indecent" material over the Internet. The CDA prohibited Internet users from using the Internet to communicate material that, under contemporary community standards, would be deemed patently offensive to minors under the age of 18. In holding that the CDA violated the First Amendment, the Supreme Court explained that without defining key terms, the statute was unconstitutionally vague. Moreover, the Court noted that the breadth of the CDA was "wholly unprecedented" in that, for example, it was "not limited to commercial speech or commercial entities" but rather embraced "all nonprofit entities and individuals posting indecent messages or displaying them on their own computers."

In response to the Supreme Court's invalidation of the CDA, Congress in 1998 enacted the *Child Online Protection Act*[144] (COPA). The COPA scaled down the restrictions of the CDA by using the "harmful to minors" category of speech. The COPA was found unconstitutional by a district court because it failed the strict scrutiny test for content-based regulation of speech.[145] That holding was affirmed on appeal, although on a different rationale—that because "harmful to minors," like obscenity, must be defined by reference to discrete, diverse, geographically defined communities,

COPA's attempt to promulgate a national "harmful to minors" standard was necessarily unconstitutional.[146] States also have attempted to control access to Internet speech by using the "harmful to minors" concept, but these attempts have proven similarly unsuccessful.[147]

The second category is exemplified by the *Child Internet Protection Act* (CIPA),[148] which, roughly speaking, requires all schools and libraries that receive certain types of federal funding to adopt and implement an Internet safety policy and to use "technology protection measures" on all computers that offer Internet access. Such measures must block or filter Internet access to visual depictions that are obscene, child pornography, or harmful to minors. Libraries must also regulate access by minors to "inappropriate material." Thus, the CIPA effectively requires the use of filtering software, sometimes referred to as *censorware*.

72.2.6 Attempts to Control Access: Case Law. The American Library Association has taken a firm stance against censorware in libraries[149] and has committed itself to challenging federal censorware requirements. To date, though CIPA applies to both libraries and schools, the legal battle over the use of censorware to control access to speech on the Internet has primarily involved libraries. The plaintiffs in these cases have library and librarian associations, including the American Library Association, library patrons and users, and entities and individuals whose Websites are likely to be blocked by censorware.

72.2.6.1 Public Libraries. On one hand, public libraries must be concerned that, by implementing censorware, they violate patrons' First Amendment right to receive information. On the other hand, public libraries also face liability for encouraging a hostile workplace for their employees by not implementing censorware on library Internet terminals.

As to the first concern, libraries face a conflict between not censoring adults' access to constitutionally protected material on the Internet and protecting children from inappropriate content. In a 1998 case involving a Virginia library, a federal district court found the library's blocking policy unconstitutional.[150] The challenged policy required that "site-blocking software … be installed on all [library computers]" so as to block child pornography, obscene material, and harmful-to-minors material.[151] The defendants implemented the policy by installing the commercial filtering software X-Stop on all public library terminals in the county.

The court first concluded that the public libraries were limited public forums and then subjected the defendants' indisputably content-based filtering policy to strict constitutional scrutiny.[152] Assuming that the library's interests were compelling, the court found that its policy was not the least restrictive means of furthering those interests. First, the court held that the library had not even considered less restrictive means, including privacy screens, library staff monitoring of Internet use (which had been used to enforce other library policies), and filtering software installed only on Internet terminals used by minors or only when minors are using them.

More important, the court found that the censorware-based policy violated the First Amendment because censorware is both overbroad and overinclusive. On one hand, the X-Stop censorware blocked more speech than that targeted by the policy.[153] On the other hand, even if censorware could somehow be tailored to exclude only restricted materials, it unconstitutionally "limit[s] the access of all patrons, adult and juvenile, to material deemed fit for juveniles."[154]

Finally, the court found that the library policy violated the doctrine barring prior restraints. Mandatory filtering policies that rely on commercial blocking software arguably constitute prior restraints because, as in *Mainstream Loudoun*, they "entrust all … blocking decisions … to a private vendor" whose standards and practices cannot be monitored by the filtering library.[155]

In a notable departure from this line of argument, the U.S. Supreme Court found in a 2003 case from Pennsylvania that CIPA was constitutional because the library's interest in preventing minors from accessing obscene material "outweighs the free speech rights of library patrons and website publishers."[156]

As to the second concern, libraries are unlikely to face liability from patrons for failing to use censorware. For example, in *Kathleen R. v. City of Livermore*,[157] a woman sued a public library, complaining that her 12-year-old son was able to view and download pornography at a public library in Livermore, California. The plaintiff sought an injunction that would prohibit the library "from maintaining any computer system on which it allows people to access … obscene material or on which it allows minors to access … sexual material harmful to minors." The plaintiff also claimed that the library had a constitutional duty to "protect" library patrons from unwanted and "harmful" sexually explicit materials. The California court of appeals rejected all of the plaintiff's claims.

A different concern here is library liability to employees for failing to use censorware under sexual harassment *hostile work environment* theories. Federal civil rights laws and many states' parallel civil rights laws afford employees the "right to work in an environment free from discriminatory intimidation, ridicule and insult."[158] Arguably, access to materials over the Internet that are offensive due to their sexually explicit nature or messages regarding race, religion, or ethnicity may subject a library to liability for a hostile environment under these laws. In 2001, employees of a Minnesota library brought a successful complaint before the Equal Employment Opportunity Commission claiming that the library failed to prevent a hostile work environment by failing to prevent the viewing of objectionable material on library computers.[159]

Such complaints and resulting lawsuits are common in many workplaces. For example,

- In 2009, a Utica, NY, woman was awarded $580,000 in damages because she was harassed and fired after she began an "interracial" relationship with a colleague while working as a security hospital treatment assistant at the Central New York Psychiatric Center:

  > Bergerson's complaint also listed several other harassing behaviors that allegedly took place, such as the use of "sexually oriented" posters and inappropriate computer screensavers in the work area, and said male coworkers told her that women were not wanted there and started rumors regarding her sex life.[160]

- A posting on the *Illinois Sexual Harassment Attorney Blog* in May 2013 argued that offensive text messages may be a form of harassment that is actionable:

  > So your boss keeps sending you sexual type text messages. He tells you how sexy you look and what he would like to do to you after work. Is this a hostile work environment in Chicago? The answer is yes. In fact this would be considered sexual harassment and you would have a claim against both the supervisor and the company. Under state law, there is strict liability against the company if a supervisor engages in sexual harassment. That means there is no requirement that the employee report the conduct before liability attaches.[161]

72.2.6.2 Public Schools. There are no cases involving the use of censorware in public schools. Schools traditionally have been considered places where the government has a strong interest in regulating speech,[162] particularly with respect to curricular materials.[163] Prevailing cases emphasize the importance of inculcating values and maintaining civility in schools, along with the risk that the viewpoints expressed by student speech would be attributed to the school. Moreover, most public school students are too young to enjoy the same constitutional rights as adults.

Nevertheless, the First Amendment should place significant limits on public schools' use of censorware. In general, students do not "shed their constitutional rights to freedom of speech or expression at the schoolhouse gate."[164] Thus, while schools have maximum authority over curriculum and school-sponsored activities and substantial authority to regulate student speech, they have less authority to regulate students' right to receive information. Despite schools' "comprehensive authority" to prescribe and control student conduct, "students may not be regarded as closed-circuit recipients of only that which the State chooses to communicate."[165]

Moreover, the extent of governmental power to regulate conduct of minors not constitutionally regulable when committed by adults is a "vexing" question, "perhaps not susceptible of precise answer."[166] Although the government has a general interest in regulating schools, the interest in regulating students' access to "objectionable" extracurricular information is of uncertain pedigree. For instance, the Supreme Court has characterized the government interest in protecting children against "harmful to minors" material as derivative of or secondary to parents' interests.[167] Thus, "[s]peech that is neither obscene as to youths nor subject to some other legitimate proscription cannot be suppressed solely to protect the young from ideas or images that a legislative body thinks unsuitable for them."[168] This characterization suggests that the state has no interest in blocking a minor's access to HTM material to which parents do not object.[169]

Moreover, the state's interest is not uniformly strong across all minors. In general, minors' fundamental rights, including the right to receive information, strengthen as they grow older; "constitutional rights do not mature and come into being magically only when one attains the state-defined age of majority."[170] Moreover, some high school students are 18 or 19 and thus not minors for HTM purposes.

The Electronic Freedom Foundation published a report in 2003 that examined the issue of Internet content filtering in public schools. Based on the state of current filtering technology, which significantly underblocks material that would be considered harmful to minors and significantly overblocks some material that would be considered protected by the First Amendment, the EFF recommended that schools should use other, less restrictive methods.[171]

In 2006, the EFF and the ALA strongly opposed a bill passed by the House of Representatives, the *Deleting Online Predators Act*[172] (DOPA). The bill would have blocked federal funding for IT in any school allowing the use of "commercial social-networking sites" such as Facebook or Myspace. The bill died in the Senate.[173]

72.3 PARENTAL INVOLVEMENT/RESPONSIBILITY. Although the technical issues of censorware are identical regardless of whether it is being applied by government, employers, or parents, the parent–child case generates much of the political discussion. Two conflicting mental models are at work in the debate.

The primary impetus for censorware arises out of an idea that might be called the "toxic material" theory, under which viewing certain information has a toxic effect,

akin to a harmful drug or a poison. As was stated in a lawsuit seeking an injunction against a library with unrestricted Internet access: "Children such as Brandon P. who view obscenity and pornography on the library's computers can and have sustained emotional and psychological damage in addition to damage to their nervous systems. It is highly likely that such damage will occur given the library's policy."[174] The outcome of the toxic material threat model has been legislation such as the CIPA.

In contrast, theoretical discussion typically takes place in a model that can be termed the "control-rights" theory. This is concerned primarily with authority relationships within some framework. One paper submitted to a government commission stated:

> [T]he decision by a third party that a person may not use a computer to access certain content from the Internet demands some sort of justification. The burden should be on the filterer to justify the denial of another person's access. The most plausible justifications for restricting access are that the third party owns the computer or that the third party has a relation of legitimate authority over the user.[175]

However, the control-rights discussion does not address the assumptions and threat model inherent in the toxic material theory. A belief in harmful information logically requires restrictions to be as extensive and widespread as possible. In fact, this is an instance of typical security practices, where the potential attack is viewed along the lines of dangerous contamination. All possible environments where the subject may be located then need to be secured to the greatest extent possible. This social need is similar to the technical requirement regarding banning privacy, anonymity, and even language-translation Websites. It is all a matter of security holes.

72.4 SUMMARY. Although the Internet has the technological potential to create a global forum for free expression, one must not underestimate the political and technological barriers at work.

Internet-specific laws. Some governments have criminalized certain types of Internet speech. Such criminal penalties may come in the form of laws intended to protect minors from "harmful" material. Under U.S. law, however, requiring Internet speakers to shield certain populations from their speech has been found to effect a ban on that speech.[176]

Application of existing laws. Governments need not specifically enact laws targeting Internet-based speech. For example, the German government action against CompuServe for providing access to illegal material merely applied existing laws to the Internet.

Content-based license terms applied to Internet users and service providers. Some countries have established licensing systems that require Internet users and/or service providers to agree to refrain from certain kinds of speech or that block access to speech as a condition of having, or using, or providing access to the Internet. For instance, China has issued rules requiring anyone with Internet access to refrain from proscribed speech.

Compelled use of computerized censorship, rating, or content labeling tools. Various techniques can prevent individuals from using the Internet to exchange information on topics that may be controversial or unpopular, may enable certain governments to enforce their own ratings systems, and may (through overblocking by overly broad filtering techniques) inappropriately block access to significant amount of nonobjectionable content.

72.5 FURTHER READING

Brown, I., and C. T. Marsden. *Regulating Code: Good Governance and Better Regulation in the Information Age*. MIT Press, 2013.

Hill, S. *Digital Revolutions: Activism in the Internet Age*. World Changing, 2013.

Howard, P. N., and M. H. Hussain. *Democracy's Fourth Wave? Digital Media and the Arab Spring*. Oxford University Press, 2013.

Kiethman, V. *Banned Websites*. Ox Books Publishing (OxBo), 2013.

Kolbert, K., and Z. Mettger. *Justice Talking From NPR (National Public Radio): Censoring the Web—Leading Advocates Debate Today's Most Controversial Issues*. The New Press, 2002.

Lessig, Lawrence. *Code and Other Laws of Cyberspace*. New York: Basic Books, 1999.

Pager, S. A., and A. Candeub, eds. *Transnational Culture in the Internet Age*. Edward Elgar Publishers, 2012.

Reed, C. *Making Laws for Cyberspace*. Oxford University Press, 2012.

Russomano, J. *Defending the First: Commentary on First Amendment Issues and Cases*. Routledge, 2013.

Youngs, G., editor. *Digital World: Connectivity, Creativity and Rights.* Routledge, 2013.

Ziccardi, G. *Resistance, Liberation Technology and Human Rights in the Digital Age*. Springer, 2012.

72.6 NOTES

1. Universal Declaration of Human Rights, art. 19, G.A. Res 217A, U.N. GAOR, 3d Sess., U.N. Doc. A/810 (1948).
2. Thomas D. Jones, "Human Rights: Freedom of Expression and Group Defamation Under British, Canadian, Indian, Nigerian and United States Law—A Comparative Analysis," 18 *Suffolk Transnat'l L. Rev*. 427, 428 (1995) (footnotes omitted).
3. "Mobile subscribers worldwide." *Global Mobile Statistics 2013 Part A: Mobile Subscribers; Handset Market Share; Mobile Operators.* mobiThinking (2013). http://mobithinking.com/mobile-marketing-tools/latest-mobile-stats/a#subscribers
4. Samuel Fifer and Michael Sachs, "The Price of International Free Speech: Nations Deal With Defamation on the Internet," 8 *DePaul-LCA J. Art & Ent. L.* 1, 2 (1997).
5. Samuel Fifer and Michael Sachs, "The Price of International Free Speech: Nations Deal With Defamation on the Internet."
6. Order of the County Court of Paris, at 2, in *La Ligue Contre Le Racisme Et L'Antisemitisme v. Yahoo! Inc.*, No. RG: 00/05308 (Nov. 20, 2000) (translated).
7. Order of the County Court of Paris, at 4.
8. Order of the County Court of Paris, at 4.
9. *Yahoo! Inc. v. La Ligue Contre le Racisme et L'Antisemitisme, et al.,* 145 F. Supp. 2d 1168 (N.D. Cal. 2001). www.cdt.org/jurisdiction/010607yahoo.pdf; Retrieved from WayBackMachine: http://web.archive.org/web/20011001001826/http://www.cdt.org/jurisdiction/010607yahoo.pdf
10. *Yahoo! Inc. v. La Ligue Contre le Racisme et L'Antisemitisme, et al.*
11. See note 4 generally.

12. Jane Ginsburg, "Copyright Without Borders? Choice of Forum and Choice of Law for Copyright Infringement in Cyberspace," 15 *Cardozo Arts & Ent. L.J*. 153 (1997).
13. Peter Swire, "Of Elephants, Mice, and Privacy: International Choice of Law and the Internet," 32 *Int'l Law*. 991 (1998).
14. *See* text at Section 72.2.5.
15. *See* text at Section 72.1.3.3.2.
16. John T. Delacourt, "Recent Development: The International Impact of Internet Regulation," 38 *Harv. Int'l L. J*. 207, 212 (1997).
17. Kim Rappaport, Note, "In the Wake of Reno v. ACLU: The Continued Struggle in Western Constitutional Democracies with Internet Censorship and Freedom of Speech Online," 13 *Am. U. Int'l L. Rev*. 765, 791 (1998).
18. Kim Rappaport, Note, at 795–799.
19. Some observers reported that the Internet seems to foster "hate" sites. Laura Leets, "Responses to Internet Hate Sites: Is Speech Too Free in Cyberspace?" 6 *Comm. L. & Pol'y* 287, 288 (2001) (noting that watchdog groups have documented about 2,800 hate sites).
20. Laura Leets, "Responses to Internet Hate Sites," at 291–294.
21. Kim Rappaport, Note, at 785–786 (footnotes omitted).
22. Kim Rappaport, Note, at 789 (footnotes omitted).
23. Steve Kettmann, "German Hate Law: No Denying It," December 15, 2000, www.wired.com/politics/law/news/2000/12/40669
24. Kettmann, "German Hate Law."
25. Kettmann, "German Hate Law."
26. J. Zittrain and B. Edelman, "Localized Google Search Result Exclusions," 2002, http://cyber.law.harvard.edu/filtering/google
27. Canadian Criminal Code §319(2), R.S.C. 1985, c. C-46 (willful promotion of hatred against an identifiable group); see *id*. at §181 (willful publication of false statements likely to injure a public interest).
28. *R. v. Keegstra*, 3 S.C.R. 697 (1990).
29. *R. v. Keegstra*, 3 S.C.R. 697 (1990).
30. Robert Cribb, "Canadian Net Hate Debate Flares," March 25, 1998, www.wired.com/news/news/politics/story/11195.html; Retrieved from *WayBack-Machine*, http://web.archive.org/web/19990203162438/http://www.wired.com/news/news/politics/story/11195.html
31. R. Howard, "Notorious Internet Service Closes; BC to continue probe of Klatt." *The Globe and Mail* (1998-04-28), www.nizkor.org/hweb/orgs/canadian/bc/fairview-technologies/press/globe-and-mail-980428.html
32. Zündel was charged with spreading "false news" in violation of §181 of the Canadian Criminal Code.
33. *R. v. Zündel* [1992] 2 S.C.R. (Can.) 731.
34. Credence Fogo-Schensul, "More Than a River in Egypt: Holocaust Denial, the Internet, and International Freedom of Expression Norms," 33 *Gonz. L. Rev*. 271, 245 n.32 (1997–1998) (citation omitted).
35. "Ernst Zündel Sentenced to 5 Years in Prison for Holocaust Denial," *Winnipeg Free Press*, February 15, 2007, online edition, www.winnipegfreepress.com/

breakingnews/world/story/3881650p-4489229c.html; Retrieved from *WayBack Machine*, http://web.archive.org/web/20070927194246/http://www.winnipegfreepress.com/breakingnews/world/story/3881650p-4489229c.html

36. Internet World Stats, March 24, 2007, www.internetworldstats.com/asia/cn.htm
37. "China Begins Push for Internet Censorship," *BizAsia*, August 23, 2000. Retrieved from *WayBackMachine*, http://web.archive.org/web/20000823190350/http://www.bizasia.com/gen/articles/stand_art.htm
38. Managing Internet Information-Release Services, P.R.C. Ministry of Information Industry Regulation, November 7, 2000; see also "China Issues Regulations on Managing Internet Information-Release Services," *China Online*, November 13, 2000. Retrieved from *WayBackMachine*, http://web.archive.org/web/20010628140620/http://www.chinaonline.com/issues/internet_policy/NewsArchive/Secure/2000/November/C00110604.asp. Other restrictions target a variety of disfavored groups, particularly supporters of the Falun Gong spiritual movement. See "China Passes Internet Security Law," China Online, December 29, 2000, www.chinaonline.com/issues/Internet_policy/NewsArchive/Secure/2000/December/C00122805.asp (URL inactive).
39. Geremie R. Barme and Sang Ye, "The Great Firewall of China," *Wired*, June 1997, www.wired.com/wired/archive/5.06/china_pr.html
40. Geremie R. Barme and Sang Ye, "The Great Firewall of China."
41. Digital Freedom Network, "Attacks on the Internet in China," April 30, 2001, www.dfn.org/focus/china/netattack.htm; Retrieved from *WayBackMachine*, http://web.archive.org/web/20010506184353/http://www.dfn.org/focus/china/netattack.htm
42. Michael Liedtke, "Google Agrees to Censor Results in China," *Associated Press Business Wire*, January 24, 2006, via AnandTech Forums, http://forums.anandtech.com/archive/index.php/t-1795022.html
43. M. Murphy, "China's State Internet Information Office Flexes Its Muscles," *Bloomberg Law*, 2011, http://about.bloomberglaw.com/practitioner-contributions/chinas-state
44. B. Blanchard and S. Huang, "China Tightens Internet Controls, Legalizes Post Deletion," *Reuters*, December 28, 2012, www.reuters.com/article/2012/12/28/net-us-china-internet-idUSBRE8BO01320121228
45. "Internet Censorship in China." *The New York Times*, http://topics.nytimes.com/topics/news/international/countriesandterritories/china/internet_censorship/index.html
46. Human Rights Watch, "The Internet in the Middle East and North Africa: Free Expression and Censorship," July 1, 1999, www.hrw.org/reports/1999/07/01/internet-middle-east-and-north-africa-free-expression-and-censorship
47. Saudi Internet regulations, Saudi Arabia Council of Ministers Resolution, February 12, 2001, www.al-bab.com/media/docs/saudi.htm
48. Human Rights Watch World Report 1999, "Freedom of Expression on the Internet," No longer available and not stored in *WayBackMachine*.
49. "Turkey Revokes YouTube Ban" *The Age.com*, March 10, 2007, www.theage.com.au/news/Technology/Turkey-revokes-YouTube-ban/2007/03/10/1173167025391.html

50. Freedom House, "New Report: Governments Grow Increasingly Repressive Online, Activists Fight Back," 2012, www.freedomhouse.org/article/new-report-governments-grow-increasingly-repressive-online-activists-fight-back
51. Freedom House, *Freedom on the Net 2012: Summary of Findings*, www.freedomhouse.org/sites/default/files/inline_images/FOTN%202012%20FINAL.pdf
52. Christina Fink, "Burma: Constructive Engagement in Cyberspace?" *Cultural Survival Quarterly* 21 no. 4 (Winter 1997), www.culturalsurvival.org/ourpublications/csq/article/burma-constructive-engagement-cyberspace
53. Digital Freedom Network, "The New Net Regulations in Burma," January 31, 2000, www.dfn.org/voices/burma/webregulations.htm; Retrieved from *WayBackMachine*, http://web.archive.org/web/20030701231239/http://www.dfn.org/voices/burma/webregulations.htm
54. Digital Freedom Network, "The New Net Regulations in Burma."
55. Alison Hunter, "Burma's Internet Censorship Worsening," October 12, 2005, www.mizzima.com/mizzima/archives/news-in-2005/News-in-Oct/12-Oct-05-40.htm
56. Digital Freedom Network, "Burma Wards Off the Internet Revolution," January 31, 2000; www.dfn.org/focus/burma/Web-crackdown.htm; Retrieved from *WayBackMachine*, http://web.archive.org/web/20011030171711/http://www.dfn.org/focus/burma/web-crackdown.htm
57. Freedom House, *Freedom on the Net 2012: Summary of Findings*.
58. Freedom House, *Freedom on the Net 2012: Summary of Findings*.
59. "In Russia, Internet Censorship 'on a Whim.'" Committee to Protect Journalists Website, June 26, 2013, http://cpj.org/2013/06/in-russia-internet-censorship-on-a-whim.php
60. The notion of "unprotected" speech is discussed in Section 72.2.4.3.
61. In the context of the right to speak, the requirement of state action is somewhat relaxed compared to other constitutional rights. For example, when private persons—not the government—use the courts against another private person's speech, "state action" is usually found. *E.g.*, *New York Times Co. v. Sullivan*, 376 US 254 (1964) (defamation); *Hustler Magazine, Inc. v. Falwell*, 485 US 46 (1988) (intentional infliction of emotional distress).
62. *Texas v. Johnson*, 491 US 397, 403 (1989).
63. *W. Va. Bd. of Educ. v. Barnette*, 319 US 624, 632 (1943).
64. *McIntyre v. Ohio Elections Comm'n*, 514 US 334 (1995).
65. *Lamont v. Postmaster General*, 381 US 301, 307 (1965) (invalidating law requiring willing recipient to request that certain, state-defined materials be sent to him); *Virginia Pharmacy Bd. v. Virginia Consumer Council*, 425 US 748, 756 (1974).
66. *City of Lakewood v. Plain Dealer Publishing Co.*, 486 US 750 (1988).
67. *Buckley v. Valeo*, 424 US 1, 19 (1976).
68. *Simon & Schuster, Inc. v. Members of N.Y. State Crime Victims Bd.*, 502 US 105, 117 (1991).
69. *City of Ladue v. Gilleo*, 512 US 43 (1994).
70. *Abrams v. United States*, 250 US 616, 630 (1919) (Holmes, J., dissenting).
71. *Hurley v. GLIB*, 515 US 557 (1995).

72. *Cohen v. California*, 403 US 15, 24 (1971) (emphasis added).
73. Bantam Books, Inc. v. Sullivan, 372 US 58, 66 (1963).
74. *Thornhill v. Alabama*, 310 US 88, 97 (1940). ("It is not merely the sporadic abuse of power by the censor but the pervasive threat inherent in its very existence that constitutes the danger to freedom of discussion.")
75. *Schenck v. United States*, 249 US 47, 52 (1919).
76. *Whitney v. California*, 274 US 357, 377 (1926) (Brandeis, J., concurring).
77. *Brandenburg v. Ohio*, 395 US 444, 447 (1969).
78. *American Booksellers Assn., Inc. v. Hudnut*, 771 F.2d 323, 333 (7th Cir. 1985), aff'd. mem. 475 US 1001 (1986).
79. See also *Planned Parenthood v. American Coalition of Life*, 244 F.3d 1007 (9th Cir. 2001).
80. *New York Times Co. v. Sullivan*, 376 US 254, 292 (1964).
81. *Police Dept. of Chicago v. Mosley*, 408 U. S. 92, 96 (1972).
82. *NAACP v. Button*, 371 US 415, 433 (1963).
83. *City of Lakewood v. Plain Dealer Publishing Co.*, 486 US 750 (1988).
84. *Smith v. California*, 361 US 147 (1959).
85. *Speiser v. Randall*, 357 US 513 (1958).
86. *Gertz v. Robert Welch, Inc.*, 418 US 323 (1974).
87. *Organization for a Better Austin v. Keefe*, 402 US 415, 419 (1971).
88. *New York Times Co. v. United States*, 403 US 713, 729 (1971) (Stewart, J., concurring).
89. *Nebraska Press Ass'n v. Stuart*, 427 US 539 (1976).
90. *Staub v. Baxley*, 355 US 313, 322 (1958).
91. *Niemotko v. Maryland*, 340 US 268, 271 (1951).
92. *Freedman v. Maryland*, 380 US 51, 58 (1965).
93. *Freedman v. Maryland* at 58–59.
94. *See, e.g., United States v. Playboy Entertainment Group, Inc.*, 120 S. Ct. 1878, 1880 (2000); *Reno v. ACLU*, 521 US 844 (1997).
95. *R.A.V. v. St. Paul*, 505 U. S. 377, 391 (1992).
96. *Perry Ed. Assn. v. Perry Local Educators' Assn.*, 460 U. S. 37, 46 (1983).
97. *R.A.V., supra* (government may not regulate otherwise unprotected "fighting words" based on viewpoint).
98. *Turner Broadcasting System, Inc. v. FCC*, 512 US 622, 642 (1994) (citations omitted).
99. *Turner Broadcasting System, Inc. v. FCC* at 661.
100. *Clark v. Community for Creative Non-Violence*, 468 US 288, 293 (1984).
101. *United States v. O'Brien*, 391 US 367, 376–77 (1968).
102. *City of Renton v. Playtime Theatres, Inc.*, 475 US 41, 47 (1986).
103. *Boos v. Barry*, 485 US 312, 335 (1988) (Brennan, J., concurring).
104. *Forsyth County v. Nationalist Movement*, 505 US 123, 135 (1992).
105. *United States v. Reese*, 92 US 214, 221 (1876).
106. *NAACP v. Button*, 371 US at 337.
107. *Kolender v. Lawson*, 461 US 352, 357–358 (1983).

108. *Thornhill v. Alabama*, 310 US 88, 97 (1940).
109. *Thornhill v. Alabama* at 97–98.
110. Often, laws are struck down only "as applied," that is, the particular application of the law to a particular person is found to violate the Constitution. Facial invalidation, on the other hand, is not limited to the particular application but to the entire law, roughly speaking. *See, e.g.*, Richard H. Fallon, Jr., "As-Applied and Facial Challenges and Third-Party Standing," 113 *Harv. L. Rev.* 1321 (2000).
111. *Broadrick v. Oklahoma*, 413 US 601, 612–615 (1973).
112. *Red Lion Broadcasting Co. v. FCC*, 395 US 367 (1969).
113. *Int'l Soc'y for Krishna Consciousness, Inc. v. Lee*, 505 US 672, 678 (1992).
114. *Legal Services Corporation v. Velazquez*, 531 US 533, (2001).
115. *Arcara v. Cloud Books, Inc.*, 478 US 697, 707 n.4 (1986).
116. *Giboney v. Empire Storage & Ice Co.*, 336 US 490, 502 (1949).
117. *Watts v. United States*, 394 US 705, 707 (1969).
118. *Watts v. United States*, 394 US at 707–8.
119. *Bartnicki v. Vopper*, 121 S.Ct. 1753, 1765 (2001).
120. *New York Times Co. v. Sullivan*, 376 US 254 (1964).
121. *Chaplinsky v. New Hampshire*, 315 US 568, 571 (1942).
122. *Chaplinsky v. New Hampshire*, at 572.
123. *Cohen v. California*, 403 US 15, 20 (1971).
124. *Miller v. California*, 413 US 15, 24 (1973).
125. *Miller v. California*, 413 US at 32.
126. *New York v. Ferber*, 458 US 747, 764 (1982).
127. The federal child pornography statute reaches "any visual depiction" of a minor under 18 years old engaging in "sexually explicit conduct," which includes "actual or simulated" sexual intercourse, bestiality, masturbation, sadistic or masochistic abuse, or "lascivious exhibition of the genitals or pubic area." 18 USC. §2256.
128. *Ferber*, 458 US at 764–765. Notably, the rationale for regulating child pornography is not content-based. Instead, it is based on the harm to children from being subjects of sexual performances, as well as the harm from the distribution of such photographs.
129. *Ashcroft v. The Free Speech Coalition*, 122 S. Ct. 1389 (2002).
130. *R.A.V. v. City of St. Paul*, 505 US 377, 391 (1992).
131. *Ginsberg v. New York*, 390 US 629 (1968).
132. *Liquormart, Inc. v. Rhode Island*, 517 US 484 (1996).
133. *Red Lion Broadcasting Co. v. FCC*, 395 US 367 (1969).
134. Speech on the Internet, however, is entitled to "full" First Amendment protection.
135. *City of Erie v. Pap's A.M.*, 120 S.Ct. 1382 (2000).
136. *FCC v. Pacifica*, 438 US 726 (1978).
137. *Denver Area Educ. Telecom. Consortium v. FCC*, 518 US 727 (1996); but *see* note 94.
138. *Sable Communications v. FCC*, 492 US 115 (1989).
139. Speech is "harmful to minors" if it (i) is "patently offensive to prevailing standards in the adult community as a whole with respect to what is suitable … for

minors"; (ii) appeals to the prurient interest of minors; and (iii) is "utterly without redeeming social importance for minors." *Ginsberg v. New York*, 390 US 629, 633 (1968) (upholding conviction of magazine vendor for selling adult magazine to 16-year-old).

140. *Reno*, 521 US at 865.
141. *Reno*, 521 US.
142. *Erznoznik v. City of Jacksonville*, 422 US 205, 212–214 (1975).
143. *Reno*, 521 US at 878. Lower courts have held that "if a work is found to have serious literary, artistic, political or scientific value for a legitimate minority of normal, older adolescents, then it cannot be said to lack such value for the entire class of juveniles taken as a whole." *American Booksellers Ass'n v. Webb*, 919 F.2d 1493, 1504–5 (11th Cir. 1990) (quoting *American Booksellers Ass'n v. Virginia*, 882 F.2d 125, 127 (4th Cir. 1989) (other citations omitted).
144. Pub. L. No. 105-277, 112 Stat. 2681 (1998) (codified at 47 USC. §231).
145. *ACLU v. Reno*, 31 F. Supp. 2d 473 (E.D. Pa. 1999).
146. *ACLU v. Reno*, 217 F.3d 162, 177 (3d Cir. 2000), cert. granted sub. nom. *Ashcroft v. ACLU*, [121 S.Ct. 1997] (2001), argued Nov. 28, 2001.
147. *E.g., ACLU v. Johnson*, 194 F.3d 1149 (10th Cir. 1999) (affirming injunction against enforcement of New Mexico statute criminalizing dissemination by computer of material harmful to minors); *American Libraries Ass'n v. Pataki*, 969 F. Supp. 160 (S.D. N.Y. 1997) (enjoining a New York statute similar to the CDA that criminalized the use of a computer to disseminate sexually explicit materials to minors).
148. Codified at 47 USC. §254(h) and 20 USC. §9134.
149. "The ALA's Resolution on the Use of Filtering Software in Libraries" is available from the ALA Office for Intellectual Freedom.
150. *Mainstream Loudoun v. Bd. of Trustees of the Loudoun County Library*, 2 F. Supp. 2d 783 (E.D. Va.1998) (*Mainstream Loudoun I*); *Mainstream Loudoun v. Bd. of Trustees of the Loudoun County Library*, 24 F. Supp. 2d 552 (E.D. Va. 1998) (*Mainstream Loudoun II*).
151. *Mainstream Loudoun I*, 2 F. Supp. 2d at 787.
152. *Mainstream Loudoun II*, 24 F. Supp. 2d at 564–565.
153. *Mainstream Loudoun II*, 24 F. Supp. 2d at 556 ("undisputed" that "sites that do not contain any material that is prohibited by the Policy" were blocked).
154. *Mainstream Loudoun II*, 24 F. Supp. 2d at 567.
155. *Mainstream Loudoun II*, 24 F. Supp. 2d at 569.
156. Electronic Freedom Foundation, "Supreme Court Supports Library Internet Blocking Law: Damages Free Speech of Library Patrons and Web Publishers," EFF.org Website, June 23, 2003, www.eff.org/Censorship/Censorware/20030623_eff_cipapr.php; Retrieved from *WayBackMachine*: http://web.archive.org/web/20080907201038/http://www.eff.org/Censorship/Censorware/20030623_eff_cipapr.php
157. 87 Cal. App. 4th 684 (1st Dist. 2001).
158. *Meritor Savings Bank FSB v. Vinson*, 477 US 57, 65 (1986).
159. "EEOC Determination that Net Porn Leads to Library Hostile Working Environment," EFF.org Website, May 23, 2001, www.eff.org/Censorship/Censorware/

20010523_eeoc_determination.html; Retrieved from *WayBackMachine*: http://web.archive.org/web/20010723005939/http://www.eff.org/Censorship/Censorware/20010523_eeoc_determination.html

160. C. Potts, "Woman Awarded $580,000 in Hostile Work Environment Case," *Utica Observer Dispatch*, October 20, 2009, www.uticaod.com/news/x536357184/Woman-awarded-580-000-in-hostile-work-environment-case
161. P. M. LaSorsa, "Are Sexy Text Messages a Form of a Hostile Work Environment in Chicago?" *Illinois Sexual Harassment Attorney Blog*. May 11, 2013, http://www.illinoissexualharassmentattorneyblog.com/2013/05/are_sexy_text_messages_a_form.html
162. *Bethel School District No. 403 v. Fraser*, 478 US 675 (1986) (upholding suspension of high school student for making sexually suggestive speech at school assembly).
163. *Hazelwood Sch. Dist. v. Kuhlmeier*, 484 US 260 (1988) (upholding teacher's censorship of articles destined for newspaper prepared by journalism class).
164. *Tinker v. Des Moines Indep. Community School Dist.*, 393 US 503, 506 (1969).
165. *Tinker v. Des Moines Indep. Community School Dist.* at 511. Schools may [, however,] regulate students' First Amendment rights if the prohibited speech would "materially and substantially interfere with the requirements of appropriate discipline in the operation of the school." *Id*. at 509. This burden cannot be satisfied by "undifferentiated fear or apprehension of disturbance"; the school must show "something more than a mere desire to avoid the discomfort and unpleasantness that always accompany an unpopular viewpoint." *Id*. at 508.
166. *Carey v. Population Servs. Int'l*, 431 US 678, 692 (plurality opinion) (1977).
167. *Reno v. ACLU*, 521 US 844 (1997); *Ginsberg v. New York*, 390 US 629 (1968); *Prince v. Massachusetts*, 321 US 158, 166 (1944).
168. *Erznoznik*, 422 US at 213–214.
169. In other contexts, schools often accommodate or try to accommodate parental wishes as to the exposure of their children in school to undesired material, even when the material is part of the school's chosen curriculum. Catherine J. Ross, "An Emerging Right for Mature Minors to Receive Information," *2 U. Pa. J. Const. L.* 223, 247 n.119 (1999) ("nearly every school district in the country allows parents 'to opt their own children out of sexuality and AIDS education, as well as out of specific activities or assignments that conflict with their religious beliefs.'"), *quoting* People for the American Way, *A Right Wing and a Prayer: The Religious Right and Your Public Schools* 60 (1997). In some states, such accommodation with respect to all or specified categories of curricular materials is required. *Minn. Stat. Ann.* §126.699 (West 1994) (schools must allow parents to review curricular materials and must provide alternative materials to replace those that parents find objectionable for any reason); *Mass. Gen. Laws Ann*. ch. LXXI 32A (West 1998) (requiring that school districts notify parents of the content of any curriculum primarily involving "human sexual education or human sexuality" and afford parents a flexible way of exempting their children from such curricula upon written notice to the school); *Va. Code Ann*. §22.1-207-2 (Michie 1998) (giving public school parents the right to review "the complete family life curricula [sic], including all supplemental materials").
170. *Planned Parenthood of Cent. Mo. v. Danforth*, 428 US 52, 74 (1976) (minors' right to abortion).

171. Electronic Freedom Foundation, "Internet Blocking in Public Schools," EFF.org Website, June 23, 2003, https://www.eff.org/node/67552
172. H.R. 5319 (109th): Deleting Online Predators Act of 2006. www.govtrack.us/congress/bills/109/hr5319
173. S. O'Hear, "Deleting Online Predators Act: R.I.P." ZDNET, January 4, 2007, www.zdnet.com/blog/social/deleting-online-predators-act-r-i-p/55
174. *First Amended Complaint in Kathleen R. v. City of Livermore*, www.techlawjournal.com/courts/kathleenr/19981103.htm
175. Members of the Information Society Project at Yale Law School, "Filtering the Internet: A Best Practices Model," September 15, 1999, www.copacommission.org/papers/yale-isp.pdf
176. *Reno v. ACLU*, 521 US 844 (1997).

CHAPTER 73

EXPERT WITNESSES AND THE *DAUBERT* CHALLENGE

Chey Cobb

73.1 INTRODUCTION 73·1

73.2 DAUBERT 73·2
- 73.2.1 Expert Witnesses' Testimony 73·2
- 73.2.2 *Daubert* Challenge 73·3

73.3 WHETHER THE *DAUBERT* CHALLENGE IS APPLICABLE: REFINING *DAUBERT* 73·3
- 73.3.1 *General Electric Co. v. Joiner* 73·3
- 73.3.2 *Kumho Tire Co. v. Carmichael* 73·3

73.4 DIVIDED WE FALL? 73·4

73.5 BEING THE BEST YOU CAN BE 73·4
- 73.5.1 Prepare Your Résumé 73·4
- 73.5.2 Find Out Exactly What Testimony Is Expected 73·5
- 73.5.3 Examine the Paperwork 73·5
- 73.5.4 Start Reading 73·5
- 73.5.5 Prepare a Written Report 73·6
- 73.5.6 Ask for Pretrial Meetings 73·6
- 73.5.7 Be Professional 73·6
- 73.5.8 Accept the Oddities 73·6

73.6 SUMMARY 73·6

73.7 FURTHER READING 73·7

73.8 NOTES 73·7

73.1 INTRODUCTION. Whenever science or technology enters the courtroom, there must surely be an expert who can give clear and proper explanations of the subject matter to the judge and jury. As new sciences and technologies have emerged, the courts have had to decide if a person is, indeed, an expert and whether or not the science is real and admissible.[1]

In 1923, the United States courts began accepting scientific evidence based on a new rule. That rule used the "general acceptance" test to determine if evidence was legitimate. This test was based on the rulings in *Frye v. United States*, which declared that if a scientific practice was generally accepted among the scientific community in which it was practiced, it could be admitted in court. This has become generally referred to as the *Frye test*.[2]

In 1975, the federal government made the scientific assertions a bit stronger by issuing the Federal Rules of Evidence No. 702 (FRE 702), which states in part:

> If scientific, technical, or other specialized knowledge will assist the trier of fact to understand the evidence or to determine a fact in issue, a witness qualified as an expert by knowledge, skill, experience, training, or education, may testify.[3]

About the same time that FRE 702 was created, various states started changing laws to include the "theft of computer services" as a punishable offense, but it was not until 1984 that the Counterfeit Access Device and Computer Fraud and Abuse Law (18 U.S.C. § 1030) was created, officially making it a federal offense to play fast and loose with computer systems of any kind.[4]

Computers had been accepted as a science, crimes could be committed on or with a computer, and someone would be needed to explain computers and their workings to the courts. Enter the digital expert witness. Although it is difficult to say when the first computer expert was used as a witness in a computer crime trial, it is fair to say that such experts soon began playing an important part in the judicial process.

Since computer crimes became "official" in 1984, computer experts are often needed in trials in which the examination of a computer or computer component is crucial to proving the case. For example, in 1991, Special Agent Jim Christy of the Air Force Office of Special Investigations (OSI) became involved in a murder case on an Air Force base where the suspect had physically destroyed two floppy disks by cutting them into many pieces. Agent Christy and a deputy played with the pieces like a jigsaw puzzle until they eventually recovered 95 percent of the data on the disks. Those data were key to obtaining a conviction.[5] Following that case, computer forensics became a hot topic and a new industry soon followed.

The computer crime laws have been updated many times in the ensuing decades, evidence rules have changed, and computer forensics are considered a solid science, but it was the *Daubert v. Merrell Dow Pharmaceuticals* case in 1993 that really began to have an effect on expert witness testimony in cases where a computer or related technology is involved.[6]

73.2 DAUBERT.

In 1974, a boy name Jason Daubert was born with a severely disfigured right arm and hand. His mother had taken Bendectin, an anti-nausea medication made by Merrell Dow Pharmaceutical, to combat the morning sickness she experienced while pregnant. This particular drug had been taken by tens of millions of women over the course of more than two decades, and a considerable number of children were born with significant birth defects. Many thousands of families took Merrell Dow to court over this problem, but Daubert's case was not decided until 1993.

73.2.1 Expert Witnesses' Testimony.

Merrell Dow stated that the drug was not the problem, and worked to get the case thrown out of court. Merrell Dow presented the court an affidavit, given by a medical doctor, stating that he had reviewed many studies that indicated that the drug offered no risks in *humans* when taken during the first trimester of pregnancy. The Daubert family countered with eight affidavits from other experts that indicated there was a significant risk of birth defects in *animal* tests.

Sadly, the court did dismiss the case. In simple terms, the court felt that the Dauberts' affidavits did not compare apples to apples (i.e., animal studies were not the same as human studies). More importantly, however, the court also said that Dauberts' experts'

findings were not admissible because they were neither published nor had they been subjected to peer review.

73.2.2 *Daubert* Challenge. The *Daubert* decision was heralded by many in the legal system as one of the most important findings by the Supreme Court in decades, because it went straight to the heart of the question of expert witnesses and testimony concerning science of one type or another. Many attorneys felt that juries were not capable of understanding new sciences or able to find flaws in scientific methods. With the *Daubert* case, however, the judges become "gatekeepers" of admissible evidence and, as gatekeepers, they decide whether scientific evidence could be admissible at all, based on these considerations:

- Whether the theory or technique used by the expert has been tested or *can* be tested
- Whether the theory or technique has ever been published and subjected to peer review
- Whether there is a known or potential rate of error for the method used
- The degree of acceptance within the relevant scientific community
- Whether standards exist for the method or technique used

73.3 WHETHER THE *DAUBERT* CHALLENGE IS APPLICABLE: REFINING *DAUBERT*. It is hard to imagine any legal ruling in a U.S. court ever going unchallenged for very long, and so it went with *Daubert*. When lawyers refer to a *Daubert* challenge today, they are more likely making reference to the *Daubert* trilogy, which are the rulings in *Daubert* plus two other subsequent cases: *General Electric Co. v. Joiner*[7] and *Kumho Tire Co. v. Carmichael*.[8] In these cases, the judges made rulings that were an attempt to further define and refine the original *Daubert* rulings.

73.3.1 *General Electric Co. v. Joiner*. The 1997 Joiner case was one concerning product liability. Experts were called up to testify to the cancer-causing abilities of a certain fluid that Mr. Joiner came in contact with on a daily basis. Ultimately, the ruling made changes to the original *Daubert* ruling by adding that the *Daubert* challenge should be used not just for scientific testimony but for *all* testimony that could be described as "scientific, technical, or other specialized knowledge" under FRE 702.

Additionally, in *Joiner* it was ruled that the judge's scrutiny is not limited to methodology but can extend to the expert's analysis. This puts another strain on expert witnesses because the science, technology, and methodologies of collecting the evidence can be fine, but the interpretation of that evidence can be called into question.

73.3.2 *Kumho Tire Co. v. Carmichael*. Last in the *Daubert* trilogy is *Kumho Tire,* which was another product liability case tried in 1999. The science in question here revolved around tire failure, and *Daubert* challenges abounded. The final ruling in this case decreed that a trial judge is vested with broad authority to determine whether and to what extent the enumerated *Daubert* factors are applicable to a case.

To clarify: Technical and nontechnical testimony can be challenged, the expert witness's findings can be challenged, and, finally, the judge can decide that all, some, or none of the *Daubert* challenges apply.

Computers are sometimes easier to understand than U.S. law.

73.4 DIVIDED WE FALL? Since *Joiner* and *Kumho Tire,* there have been other decisions concerning *Daubert*, admissibility of evidence, and expert witnesses, but the *Daubert* trilogy seems to be the standard—at least for federal courts. (*Daubert* and subsequent decisions were made in federal courts.) State courts, however, seem to have a lot more leniency in deciding what rules to use to determine admissibility. If you are an expert witness in a federal trial, there is a good chance you may have to face a *Daubert* challenge. If the case is at the state level or lower, you will likely experience an unusual mix of rules and challenges. The next list shows which standards apply in different states of the United States.

States Accepting *Daubert*	States Accepting *Frye*	States with Their Own Rules
Connecticut	Alaska	Arkansas
Indiana	Arizona	Delaware
Kentucky	California	Georgia
Louisiana	Colorado	Iowa
Massachusetts	Florida	Military
New Mexico	Illinois	Minnesota
Oklahoma	Kansas	Montana
South Dakota	Maryland	North Carolina
Texas	Michigan	Oregon
West Virginia	Missouri	Utah
	Nebraska	Vermont
	New York	Wyoming
	Pennsylvania	
	Washington	

If you are called as an expert witness in any case, it is best to know what rules and challenges the court may be considering in the case. Leave no doubt; talk to your attorney and, if he or she has not considered a possible *Daubert* (or other) challenge, it could mean a tough road ahead. A competent attorney will be able to give you guidance on what testimony is likely to be challenged, and an excellent attorney will spend a long time with you going over your questions and concerns.

73.5 BEING THE BEST YOU CAN BE. Being an expert witness in a trial means more than just sitting in the stand and letting your expertise and knowledge flow. You need to be on the top of your game and know absolutely all you can possibly know about the subject matter in which you are to give testimony. One court recently dismissed an expert witness's opinion in digital forensics because he had *heard* of a certain book but had not personally read it.[9]

Although there are no hard-and-fast guidelines to being a good expert witness, you can use some general stratagems that have worked well for others. Certainly, if the list of books in online bookstores is any indication, the market for expert witnesses is strong.[10]

73.5.1 Prepare Your Résumé. The first step in becoming an expert witness is to have a very strong résumé that is completely and totally factual. The résumé will be entered in the trial records, so this is no time for elaboration of the facts. Expert

testimony is allowed only if the court declares a witness to be an expert in a specified field, and your résumé will help establish the facts of the matter.

Highlight these items in your résumé to indicate your areas of expertise:

- Education
- Training
- Certificates
- Licenses
- Employment
- Lists of published articles, books, opinions
- Lists of speaking engagements and titles of talks

Once the court has examined your résumé, oral examination by the attorneys and the judge will likely follow. This can be quite easy, or it can be long and arduous.

73.5.2 Find Out Exactly What Testimony Is Expected. When contacted by an attorney to be an expert witness, do not be too fast to accept or decline. Most likely the attorney will be fishing at this point and will discuss things only in broad, general terms. It is also likely that your contact has limited knowledge of the subject matter. Find out as much as you can about the charges and what the attorney needs (or wants) to make the case.

Once you have reached an understanding with your contact, get everything in writing and make it a contract. This will be a statement of what is expected of you and what expertise is required to give the testimony. Your contract should also list your fees and payment schedule. This is where TV and reality take different roads. You have probably seen televised courtroom dramas where the opposing counsel undermines an expert witness by asking "How much were you paid for your testimony today?" Although that makes good television, in reality the courts *expect* that an expert witness is paid, and a fee actually lends credibility to the witness.

Also, look for any legal protections you will be offered for your testimony. You might be asked to do or say something that is a potential breach of ethics or confidentiality, and this could result in a loss of certain types of certification.

73.5.3 Examine the Paperwork. As soon as you are under contract, review the paperwork. If you discover there have been misrepresentations and you no longer want the job, you can back out at this point with no hard feelings. Of course, you may want to remain on the case, but you will need to discuss any concerns or questions with the attorney and possibly make an amendment to your contract.

Go over the paperwork as many times as you can until you are completely familiar with all aspects of the case. If the trial is particularly contentious, your familiarity with the case will certainly come into question.

73.5.4 Start Reading. Now is a good time to start brushing up on your knowledge base. Dig into your library of industry standard books and give them a review. Look for recent books on the subjects and articles in reputable journals. The more you are able to quote facts and figures, the better you will look on the stand.

73.5.5 Prepare a Written Report. Whether or not a written report is required should be a condition of your contract, but it is always a good idea to write up a report, even if it is just for your own use. While on the stand, you will be allowed to refer to your notes and to your report, although the court may want to examine them. Notes prepared at the time observations were made have stronger evidentiary value than statements from memory. This is a good strategy and makes you look very professional, and strong in your opinions.

It is a good idea to list all your sources and resources in your report, as well as your methodologies. This is an area where the *Daubert* challenge can get sticky, and the veracity of your opinion, based on what you have used and how you have used it, may come into question.

73.5.6 Ask for Pretrial Meetings. It is always a good idea to practice what you need to say, especially when it is to be delivered to a large group of people. Just as you would practice giving a lecture or speech, you should practice what you are going to say on the stand. You will need to appear to be confident at trial, so ask the attorney for assistance in preparing for trial. There is always the possibility that new questions or concerns may arise during the pretrial meeting, but it is better that these come up in the meetings than in front of the court.

73.5.7 Be Professional. Everyone is familiar with the geek stereotype—long hair, torn jeans, T-shirt, and sandals—but that is totally inappropriate. Dress professionally for court. When giving your testimony, use language that the layperson can understand. Do not let your ego show, and certainly contain your temper.

73.5.8 Accept the Oddities. No one expects you to understand everything that happens in a court proceeding. Simply answer the questions directly, and do not be afraid to say you do not know. (If you do not know because it cannot be known, however, that is important to relate.) You do not have to justify anything you say unless you are asked to. If an objection is raised, stop speaking and wait for the judge to ask you to continue. Last, do not be afraid to ask questions yourself if you are unclear on questions or procedures.

73.6 SUMMARY. It is certainly an ego boost to be considered enough of an expert in your field to be called as an expert witness, and some people have made successful careers of it. If you do make a career of being an expert witness, however, you cannot justify being an expert if you do not continually upgrade your expertise through industry experience. In any case, whether you are an expert, or whether your testimony is acceptable, is ultimately up to the court.

If you are called to be an expert witness and agree to the job, you do not really need to know the legal ins and outs of the *Daubert* challenge (or any other challenge); that is the attorney's job. The best plan is to know your subject matter better than anyone else and to be completely firm in your opinions and conclusions. If there is a challenge, do not take it personally. More often than not, the challenge has more to do with the methodology used in coming to your conclusions or with the science itself. To be prepared, it is best to keep detailed notes of everything you read and every step you take in coming to your opinions.

73.7 FURTHER READING

Brodsky, S. L. *Testifying in Court: Guidelines and Maxims for the Expert Witness*, 2nd ed. American Psychological Association, 2012.

Eastman, G., V. M. Howell, and M. Salgado. "A Primer on When to Use Expert Witnesses and How to Find Them." *Bloomberg Law* (2013). http://about.bloomberglaw.com/practitioner-contributions/a-primer-on-when-to-use-expert-witnesses-and-how-to-find-them/

Hallberg, B. J. *How To Be An Effective Litigation Consultant And Expert Witness*. AuthorHouse, 2010.

Lubet, S. and E. Boals. *Expert Testimony: A Guide for Expert Witnesses and the Lawyers Who Examine Them*, 2nd ed. LexisNexis/National Institute for Trial Advocacy, 2009.

Poynter, D. *The Expert Witness Handbook: Tips and Techniques for the Litigation Consultant*, 3rd ed. Para Publishing, 2007.

Sennewald, C. A. *From the Files of a Security Expert Witness*. Butterworth-Heinemann, 2013.

73.8 NOTES

1. The author is not a lawyer, and the facts presented here are a layperson's interpretation of the laws concerning digital evidence and expert witnesses.
2. *Frye v. United States,* 293 F. 1013 (D.C. Cir. 1923), www.daubertontheweb.com/frye_opinion.htm
3. Federal Rules of Evidence (amended to December 1, 2011), www.law.cornell.edu/rules/fre/rules.htm
4. Counterfeit Access Device and Computer Fraud and Abuse Act of 1984, www.answers.com/topic/counterfeit-access-device-and-computer-fraud-and-abuse-act-of-1984
5. Robin Mejia, "CSI: TCP/IP." *Wired*, Issue 15.01, January 2007, www.wired.com/wired/archive/15.01/cybercop.html
6. *Daubert v. Merrell Dow Pharmaceuticals,* 509 U. S. 579 (1993), www.law.cornell.edu/supct/html/92-102.ZS.html
7. *General Electric v. Joiner,* 522 U. S. 136 (1997), www.law.cornell.edu/supct/html/96-188.ZO.html
8. *Kuhmo Tire v. Carmichael,* 526 U. S. 137 (1999), www.griffithlaw.net/research.database/uploads/DAUBERT%20CHALLENGES%20IN%20GOVERNMENTAL%20LITIGATION%20final%20draft.pdf
9. Daniel J. Ryan and Gal Shpantzer, "Legal Aspects of Digital Forensics," n.d., http://euro.ecom.cmu.edu/program/law/08-732/Evidence/RyanShpantzer.pdf
10. Malack Amenya, "Recovering, Examining and Presenting Computer Forensic Evidence in Court," n.d., http://lib.convdocs.org/docs/index-267264.html

CHAPTER 74

PROFESSIONAL CERTIFICATION AND TRAINING IN INFORMATION ASSURANCE

M. E. Kabay, Christopher Christian, Kevin Henry, and Sondra Schneider

74.1 INCREASING COMPETENCE AND CONFIDENCE IN CYBERSECURITY PROFESSIONALS 74·1
- 74.1.1 Recognition by Government Agencies 74·2
- 74.1.2 Certification, Qualification, and Validation of Hands-on Skills 74·3
- 74.1.3 Effective Workforce Training 74·5
- 74.1.4 Career Value of Certification 74·6
- 74.1.5 Summary Chart 74·8

74.2 ACCREDITATION 74·8
- 74.2.1 ANSI/ISO/IEC 17024:2012 Accreditation of Personnel Certification 74·8
- 74.2.2 ICE/NOCA/NCCA 74·9
- 74.2.3 CNSS Standards 74·9

74.3 IA CERTIFICATIONS 74·10
- 74.3.1 ACFE 74·11
- 74.3.2 ASIS 74·11
- 74.3.3 Brainbench Basic Security Certifications 74·11
- 74.3.4 CyberSecurity Institute 74·13
- 74.3.5 CERT 74·14
- 74.3.6 CompTIA 74·14
- 74.3.7 EC-Council 74·15
- 74.3.8 CWNP 74·16
- 74.3.9 GIAC 74·16
- 74.3.10 HTCN 74·18
- 74.3.11 IACIS 74·18
- 74.3.12 IIA 74·19
- 74.3.13 ISACA 74·19
- 74.3.14 $(ISC)^2$ 74·20
- 74.3.15 mile2 74·22
- 74.3.16 Security University 74·22

74.4 PREPARING FOR SECURITY-CERTIFICATION EXAMS 74·23

74.5 CONCLUDING REMARKS 74·24

74.6 FURTHER READING 74·24

74.7 NOTES 74·25

74.1 INCREASING COMPETENCE AND CONFIDENCE IN CYBERSECURITY PROFESSIONALS. Information assurance (IA) has grown into one of the most important and sought-after fields in the job market today. This is not a market saturated with professionals. Security professionals are well paid right out of

college. And at the time of writing (August 2013), there is still a large demand for IA practitioners.

Persons responsible for a company's information security are given a lot of power. With this in mind, a company should hire only those whom it can trust to do the job right and who are ethical. Certifications provided by security groups and societies support management's quest for trustworthy, competent employees. Through rigorous examinations, adherence to approved protocols, and standards of ethics, a security professional can achieve one or many professional security-related certifications.

Perhaps one of the most critical decisions an organization has to make today is how to invest in its staff. Technology, policies, and well-defined processes are all important elements of an organization's infrastructure, but the most valuable and versatile asset that any company has today is its personnel.

How do we recruit, develop, and maintain a skilled and professional staff with other organizations competing for our trained experts, and how do we maintain their level of proficiency despite the rapid changes in technology and the many pressures of keeping the business operating?

Many companies struggle with the need to provide opportunity for their staff, while realizing that training and educating their staff makes these people attractive to other organizations.

Government regulations and legislation as well as increasing expectations from customers, suppliers, and business partners require information technology (IT) staff who are up to date with the latest technologies and who can provide assurance of reliable and stable IT operations. In most cases, a company is required either to develop internal resources to meet these demands or to pay a large price to recruit such people from external sources. More and more organizations are building relationships with universities and technical colleges to ensure that a ready supply of new graduates is available for their use. However, to retain good employees, it is necessary to make expensive investments in their further education. This is money and time that is hard to justify in some cases, and there are numerous options when it comes to registering for a professional training or educational program.

74.1.1 Recognition by Government Agencies. Upon taking office in 2009, U.S. Department of Homeland Security (DHS) Secretary Janet Napolitano recognized the urgency of the situation in cybersecurity. In October 2010 she announced a new hiring authority to enable DHS "to recruit the best cyber analysts, developers and engineers in the world to serve their country by leading the nation's defenses against cyber threats."[1]

In 2009 President Obama declared cybersecurity to be one of the nation's most serious economic and security challenges. The federal government needs a coordinated, sustained effort to build the capability of a qualified and validated cybersecurity workforce to combat cyberthreats and ensure the nation's safety.[2]

Finding the people with the needed skills is not easy. According to the DHS *Cyber-Skills Task Force Report* of Fall 2012,

> Finding the people with the needed skills, however, poses a dilemma. The numbers of professionals with these mission-critical skills are so limited that government contractors and federal agencies compete with one another and the private sector to hire them. Not surprisingly, a recent article in *Bloomberg News* reports that "the competition is fiercest for workers with hands-on experience defending systems against hackers and malicious viruses that can steal sensitive

government data."[3] For DHS to acquire sufficient talent in such a competitive environment, it needs to radically expand the national pipeline of professionals with sophisticated technical cybersecurity skills.[4]

Key recommendations from the DHS CyberSkills Task Force include the following:

- Adopt and maintain an authoritative list of mission-critical cybersecurity jobs, particularly in areas like penetration testing, security monitoring, and incident response, and modify that list in the face of changing threats and technologies.
- Develop training scenarios that allow it to properly evaluate cybersecurity talent for each of the mission-critical tasks.
- Adopt a sustainable model for assessing the competency and progress of existing and future cybersecurity talent.
- Establish a department-level infrastructure that oversees the development of the cybersecurity workforce.
- Streamline the hiring process and make government cybersecurity jobs more enticing by emphasizing the service, skills, and growth potential in the Federal Government.
- Establish a two-year, community-college-based program that identifies and trains large numbers of students for in-demand cybersecurity jobs.
- Raise the eligibility criteria for schools that participate in the Centers for Academic Excellence and Scholarship for Service programs to ensure that graduates are better prepared to perform in-depth cybersecurity work.
- Launch a major initiative to enhance the opportunities for U.S. veterans to be trained for and hired in mission-critical cybersecurity jobs.
- Use the large majority of its direct hiring authority to bring on people with critical cyber-skills, until at least 600 of those workers are fully on board.
- Specify the skills and level of proficiency needed in all cybersecurity-related contracting.
- Establish a pilot DHS CyberReserve program that ensures former DHS cybersecurity workers and other cyberprofessionals outside of government are known and available to DHS in times of need.[5]

How do we identify and capture those skilled IA professionals? How do we assess the current workforce to become the next qualified and validated cybersecurity workforce? What are the best ways to develop a national security workforce capable of meeting current and future IA challenges; how do we maintain their level of proficiency despite the rapid changes in technology and the many pressures of keeping the business operating; and, lastly, how do we improve capability to recruit and retain that sophisticated IA talent, because training makes these people attractive to other organizations?

Information-security certification has never been more important. Certification is one of the hottest topics in information security. Educating the workforce is the glue that reduces risk across the enterprise. Without hands-on education and training, daily threats and network vulnerabilities increase risk.

What are the ways to boost the nation's overall ability to develop professionals with the advanced IA skills needed to protect the information, systems, and networks that enable effective and secure operation of government and commercial elements of critical national infrastructure?

Companies need to focus the large majority of near-term efforts in IA hiring, training, and human-capital development to build teams with mission-critical IA skills.

74.1.2 Certification, Qualification, and Validation of Hands-on Skills.

The plain-English definition of a *certificate* is "document providing official evidence:

an official document that gives proof and details of something such as personal status, educational achievements, ownership, or authenticity."[6]

Certificates and certification are important in affirming the qualifications of IA professionals.

1. Any program that does not have an examination is a *certificate-of-participation* program.
2. Any examination that measures learning objectives is a part of a *certification* program.
3. A *professional certificate* program is an academic or commercial effort that awards a certificate from attending or participating in a course and *may or may not* measure the learning objectives with an examination at the course's conclusion.

Professional certificates are a way to move a career forward. Typically, certificate programs are a multiclass offering designed to expand expertise in a specific subject. They comprise required courses that cover areas of interest in more granular topics of specialization. Knowledge and hands-on skills gained can be immediately applied in the workplace.

Certification is even more valuable. Doctors take examinations to determine book or course knowledge, plus a residency to do surgery without harming anyone; pilots take courses and go through flight time to show that they can fly a plane without crashing. In IA, finding people who actually have qualified and validated cybersecurity skills is a challenge. As one commentator wrote in August 2013,

> Cybersecurity is one of the most high-demand, lucrative careers in the current IT job market. According to a report by Burning Glass Technologies, the total number of available jobs in IT security has grown by 73% from 2007 to 2012. That's 3.5 times faster than other IT disciplines and 12 times faster than the average for all other jobs… .
>
> There are a number of information security certifications, with broad and magnified focuses, that are imperative for any cybersecurity professional worth his or her salt to know. Security certifications provide a potential employer with a twofold incentive for adding you to the team. One, it shows that an impartial board of accredited professionals is willing to stake their reputations on your subject matter mastery. Two, it shows that you've been schooled in a high standard, ethical approach to your job. That way, they'll feel assured that you're not a Black Hat masquerading as a cybersecurity tech for just long enough to infiltrate their system and gorge yourself on the data held inside.[7]

Are there differences in certifications? Yes, there are hands-on skills *performance-based* certifications and *theory-based* certifications.

- *Theory-based* certification can be instructor-led lectures, and online or self-paced instruction. These programs end with an examination to measure the learning outcomes but are not intended to qualify or validate the participants' hands-on skills.
- *Performance-based* certifications are instructor-led, with daily labs, an examination to measure the learning outcomes, and written practical exams to qualify and validate that the participant has learned the intended outcome objectives.

  Today's cybersecurity workforce must become qualified and validated by a common set of cybersecurity roles/tasks/jobs standards for educating and building the cybersecurity workforce.

74.1.3 Effective Workforce Training. There is a significant difference between training and education. In this chapter, the emphasis is on training; in Chapter 75 of this *Handbook*, it is on education.

When a company invests in a technology such as firewalls or wireless systems, it often has to provide training to the people responsible for managing the devices. Such training is best provided through official sources approved of by the vendors, because product vendors can often arrange for training as a part of the purchase price; in addition, there is a certain assurance that the training materials are up to date if they are provided by the vendors or by an official source coordinating with the vendors.

74.1.3.1 Effective Training Programs. A common mistake companies make is to send one person to the official training program, trusting that person to pass the knowledge gained to their peers. What usually happens is that the person who attends the course provides a 10-minute overview or one page review to management and the rest of the staff when she returns. In theory this might work, but the knowledge transfer is limited because many organizations have no systematic way of sharing knowledge or evaluating a training program to ensuring that they receive value for their investment.

All worthwhile education and training programs should go beyond merely passing on information passively; they should provide the students with some form of practical work—whether through labs, a case study, or the need to prepare for and pass a difficult project or an extensive examination. It is too easy to have programs that are not really effective—the students can easily slip into a state of *death by PowerPoint* or *lecture hypnosis*, where they are following along with the program but not really absorbing or digesting the material. It is only when they are challenged to answer questions about the material or to use it in a practical manner that most people will truly grasp and retain the information. The goal is to integrate knowledge into the participants' world-view, not simply to pass an exam.[8]

74.1.3.2 Supporting Participants. One factor to consider is the location of the class. Ideally, a student should not be continuously distracted or pulled from the program to address issues back at the office or to handle email during the class. Distraction during class makes what is being presented or discussed in the class difficult to absorb. Lack of concentration causes the organization to lose the investment it made in the class; students may ask questions that were addressed earlier and thus frustrate other students; and distracted students may do poorly in the lab work or examinations and cause conflict by blaming the course or the instructor. While a student is assigned to a training class, she should be solely focused on the class and not participate in activities back at the office except in an emergency. Training should be held at facilities that enable serious study and learning. The courses themselves should not be just interesting hot topics, but rather those that will directly benefit in their work—and as soon as possible. Avoid educating someone in skills they will not use for several months: they will probably have forgotten most of the information by the time they try to use it.

Education and training are sometimes seen as a reward for an employee or a measure of respect for the staff. Although reward and respect are important, it is inappropriate to hand out training assignments as if they were perquisites of rank or signs of management favor. Courses should be focused on meeting specific behaviorally defined objectives.[9] In addition, to emphasize that learning is legitimate work and not a perk, employees

who take assigned courses on weekends or outside their normal work shifts should be assigned compensatory time to make up for the use of their personal time.

74.1.4 Career Value of Certification. In the *2007 SC Magazine/EC-Council Salary Survey*, a key finding was that "[w]hile not a mainstream occurrence, many companies are striving to ensure they're viewed as being proactive about securing digital assets by asking that new hires have a certification."[10] Indeed, there has been a steady increase in the acceptance of certification as a credible indicator of professional competence in the information assurance field. For example, even by 2003, the information technology industry had recognized the value of the *Certified Information Systems Professional* (CISSP) designation. A writer for the prestigious *SC Magazine* wrote:

> The CISSP designation from the International Systems Security Certification Consortium—or (ISC)^2—is often referred to as the de facto standard when it comes to information security certifications. It is described as the "International Gold Standard," according to (ISC)^2. But companies and government agencies also place importance on more technical certifications, such as the SANS Institute's Global Information Assurance Certification (GIAC).[11]

In June 2004, the CISSP designation was accredited by the International Standards Organization under the ISO/IEC 1702 standard, making it a globally recognized security-professional certification.[12]

In August 2004, the United States Department of Defense issued DoD Directive 8570.1. The DoD Fact Sheet from (ISC)^2 is provided in full as an appendix to this proposal.[13] The key points of this development are:

- "… requires any full- or part-time military service member, contractor, or foreign employee with privileged access to a DoD information system, regardless of job or occupational series, to obtain a commercial information security credential accredited by ANSI or equivalent authorized body under the ANSI/ISO/IEC 17024 Standard… ."
- The CISSP designation was recognized at the highest levels of both the technical and management categories.

The *Fact Sheet* continues with a discussion of "the significance of this mandate and of commercial certification in general." The authors wrote:

> This mandate will have far-reaching implications, including:
>
> - The Directive is viewed as a government endorsement of the effectiveness and cost-efficiency of commercial certification.
> - It provides military and civilian personnel with a certification that is professional, internationally recognized, and vendor-neutral (not tied to any agency, technology, or product).
> - It provides a portable certification that is recognized in both the public and private sectors.
> - It mandates and endorses a global standard (ANSI/ISO/IEC 17024).
> - It positions the information security profession as a distinct job series.

In a 2008 survey sponsored by (ISC)^2, analysts found that "[t]he clear differentiator in terms of career advancement and professional growth lies in certifications such as the CISSP. 'Certifications are a measure of experience and knowledge, and businesses increasingly like to see that certification… .' "[14]

Lauren Gibbons Paul wrote in April 2013,

> When it comes to education, most people agree, more is better… . Anyone who puts in the time and spends the money to get certified is showing they care about staying current with security trends and techniques. That quality makes someone more desirable to an employer… .
>
> As a practical matter, many of today's information security certifications require much hands-on application of skills… . Being able to pass proves to a potential employer that you can do certain things, potentially giving you an edge over those who do not hold the certification.
>
> For some jobs, obtaining a particular security certification, , , is a [pre]requisite for even being considered… . Beyond that, however, attaining certifications is generally a matter of personal and/or employer choice. Some certifications require a great deal of work both in and out of the classroom, as well as sitting for the test… .
>
> Whether or not security certification will earn you more, now or in the future, depends a lot on the organization, the job, and the industry. If your company values continuing education (and will help foot some of the bill for the training), that is a good indication that certification will elevate your status. If not, you may still want to pursue certification if … [for you] education is its own reward, or you need to build up your resume in anticipation of a making a move.[15]

The 2011 $(ISC)^2$ Global Information Security Workforce Study reported,

> According to the 74 percent of respondents who were involved in hiring information security staff, the importance of information security certifications as a hiring criterion remains high, with just less than 90 percent of respondents ranking security certifications very or somewhat important… .[16]

However, the authors of that *Study* warned,

> Complexity has been added to the hiring process through the years due to the sheer number and qualitative differences of certifications offered in the marketplace. The list of vendor-neutral and vendor-specific security certifications grows every year, making it difficult for employees, hiring managers, and their organizations to discern which certifications carry the greatest value for them. Today, the number has significantly grown to more than 40 vendor-neutral and vendor-specific certifications. Frost & Sullivan [the firm conducting the research for $(ISC)^2$] first reported the perception of the dilution effect in the marketplace in the 2008 study. This trend continues to challenge certification vendors to differentiate themselves. The concern is that certifications considered of high value today may be perceived to be devalued and, consequently, less significant to information security professionals and, more importantly, their employers. Frost & Sullivan believes that many vendors are addressing the current skills gap increasing the complexity and continuing education requirements of their certifications.[17]

Figure 27 from the *Study* shows that 69 percent of the 10,413 respondents rated employee competence to be a contributing reason to hire a certified employee; 54 percent reported their quality of work as a positive factor.[18]

The 2013 version of the *Workforce Study* reported,

> Slightly more than 46 percent of all survey respondents indicated that their organizations require certification, and among those respondents, 50 percent of member and 39 percent of nonmember indicate certification is a requirement. Government-defense is most emphatic on this point; 84 percent state certification is required, and a distant, but still high, second is info tech at 47 percent.[19]

The 2013 *Study* found that 68 percent of the 12,396 respondents[20] accepted employee competence and 53 percent reported quality of work as reasons for requiring information-security certifications. This consistency between 2011 and 2013 supports

EXHIBIT 74.1 Comparison of Certification and Certificates

Certification	Certificate
Results from an assessment process that recognizes an individual's knowledge, skills, and competency in a particular specialty	Results from educational process
Typically requires professional experience	For newcomers and experienced professionals
Awarded by a third-party, standard-setting organization	Awarded by educational programs or institutions often for profit
Indicates mastery/competency as measured against a defensible set of standards, usually by lab, application, and exam	Indicates completion of a course or series of courses with a specific focus (different from a degree-granting program)
Standards set through a defensible, industry-wide process (job analysis/role delineation) that results in an outline of required knowledge and skills	Course content determined by the specific provider or institution, not standardized
Typically results in credentials to be listed after one's name (e.g., CISM, CISSP, GIAC, C—EH, Q/ISP, Security+)	Usually listed on a resume detailing education
Has ongoing requirements in order to maintain; holder must demonstrate (s)he continues to meet requirements	Demonstrates knowledge of course content at the end of a set period in time

the view that certification plays a valuable role in improving job prospects for IA personnel.

74.1.5 Summary Chart. Exhibit 74.1 provides a comparison of certification and certificates.

74.2 ACCREDITATION. The next sections explain the three most important accreditation organizations and processes for information assurance: ANSI/ISO/IEC 17024, ICE/NOCA/NCCA, and IACE.

74.2.1 ANSI/ISO/IEC 17024:2012 Accreditation of Personnel Certification. The Accreditation Program for Personnel Certification Bodies under ANSI/ISO/IEC 17024, the International Organization for Standardization (ISO), and the International Electrotechnical Commission (IEC) have developed a global, voluntary benchmark for organizations responsible for certification of personnel. Fully enacted as of April 1, 2003, this international standard (ANSI/ISO/IEC 17024) was designed to harmonize personnel certification processes worldwide and create global standards for certifying the workforce in a wide range of fields. A new version (ANSI/ISO/IEC 17024:2012) was approved in July 2012.[21]

ANSI accreditation process involves both a review of a paper application and the performance of an audit (on-site visit) to validate information provided by each applicant organization with regard to the process and delivery of the examination. ANSI does not accredit/review or otherwise value the content of the certification program.

ANSI operates the Program in accordance with *ISO/IEC 17011, Conformity assessment—General requirements for accreditation bodies accrediting conformity*

assessment bodies.[22] The requirements of this document provide the basis for a nationally recognized, open, transparent, and credible accreditation activity that requires applicants' certification programs to demonstrate integrity and technical and administrative quality.

Certifications accredited by ANSI are marked with asterisks (*) below.

74.2.2 ICE/NOCA/NCCA. The *Institute for Credentialing Excellence*,[23] formerly the *National Organization for Competency Assurance* (NOCA), founded in 1977, promotes excellence in credentialing worldwide through these services:

- Education
- Research
- Advocacy
- Standards
- Accreditation

The organization provides annual conferences, Webinars, and publications. In 1987, NOCA created the *National Commission for Certifying Agencies* (NCCA) to formalize the process of accreditation of certification programs and organizations.[24]

NCCA uses a peer review process to:

- Establish accreditation standards
- Evaluate compliance with the standards
- Recognize organizations/programs that demonstrate compliance
- Serve as a resource on quality certification

Certification organizations that submit their programs for accreditation are evaluated based on the process and products, not the content, and are therefore applicable to all professions and industries.

NCCA-accredited programs certify individuals in a wide range of professions and occupations, including nurses, automotive professionals, respiratory therapists, counselors, emergency technicians, crane operators, and more. As of August 2013, NCCA had accredited 297 certification programs, mostly in the healthcare professions, but including:

- American Board for Certification in Homeland Security (ABCHS)[25]
- American Health Information Management Association (AHIMA)[26]
- Defense Security Service (DSS)[27]

74.2.3 CNSS Standards. A set of standards used in accreditation, as shown in Section 74.2.2, are defined by the *Committee on National Security Systems* (CNSS), a U.S. interagency group which describes its mission as follows:

> The CNSS provides a forum for the discussion of policy issues, and is responsible for setting national-level Information Assurance policies, directives, instructions, operational procedures, guidance, and advisories for U.S. Government (USG) departments and agencies for the security of National Security Systems (NSS) through the CNSS Issuance System. The CNSS is directed to assure the security of NSS against technical exploitation by providing: reliable and

continuing assessments of threats and vulnerabilities and implementation of effective countermeasures; a technical base within the USG to achieve this security; and support from the private sector to enhance that technical base assuring that information systems security products are available to secure NSS.[28]

The CNSS focuses on *national security systems,* which it defines as follows:

- National security systems are information systems operated by the U.S. Government, its contractors, or agents that contain classified information or that:
  - involve intelligence activities;
  - involve cryptographic activities related to national security;
  - involve command and control of military forces;
  - involve equipment that is an integral part of a weapon or weapons system(s); or
  - are critical to the direct fulfillment of military or intelligence missions (not including routine administrative and business applications).

The definition for a National Security System, along with other applicable terms used in the National Security Community, are found in CNSSI 4009, "Information Assurance Glossary."[29,30]

Some of the key standards defined by the CNSS[31] include

- NSTISSI-4011: National Training Standard for Information Systems Security (INFOSEC) Professionals, dated 20 June 1994[32]
- CNSSI-4012: National Information Assurance Training Standard for Senior Systems Managers, dated June 2004; Supersedes NSTISSI No. 4012, dated August 1997[33]
- CNSSI-4013: National Information Assurance Training Standard for System Administrators (SA), dated March 2004[34]
- CNSSI-4014: Information Assurance Training Standard for Information Systems Security Officers, dated April 2004; supersedes NSTISSI No. 4014, dated August 1997[35]
- NSTISSI-4015: National Training Standard for Systems Certifiers, dated December 2000[36]
- CNSSI-4016: National Information Assurance Training Standard for Risk Analysts, dated November 2005[37]

For further discussion of certification of academic institutions involved in IA, see Chapter 75 in this *Handbook*.

74.3 IA CERTIFICATIONS. IA has grown into one of the most important and sought-after career fields in the job market today. Security professionals are well paid, based on knowledge and experience. The demand has grown for qualified and validated IA practitioners.[38]

Persons responsible for a company's information security are given a lot of responsibility and power. An organization should hire only those whom it can trust to do the job right, and who are ethical. Certifications provided by security groups and societies support management's quest for trustworthy, competent employees. Through

rigorous examinations, adherence to approved protocols and standards of ethics, a security professional can achieve one or many professional security-related certifications.

What follows are brief summaries of some of the key security-related certifications including basic, intermediate, advanced, forensics, anti-hacking, and specialized. For those starting up the information-security career ladder or who already have security experience and wish to hone their skills in some specialized area, these resources are of great value. The sections are in alphabetical order of the certification bodies.[39]

74.3.1 ACFE. The nonprofit *Association of Certified Fraud Examiners* (ACFE)[40] offers the *Certified Fraud Examiner* (CFE) certification. The CFE demonstrates ability to detect financial fraud and other white-collar crimes. This credential is of primary interest to full-time security professionals in law or law enforcement, or those who work in organizations with legal mandates to audit for possible fraudulent or illegal transactions and activities (such as banking, securities trading, or classified operations). The ACFE notes that "According to the *2013/2014 Compensation Guide for Anti-Fraud Professionals*, CFEs earn 25 percent more than their noncertified colleagues."[41] In addition,

> To become a CFE, an individual must pass a rigorous test on the four major disciplines that comprise the fraud examination body of knowledge:
>
> - Fraud Prevention and Deterrence
> - Financial Transactions and Fraud Schemes
> - Investigation
> - Law[42]

74.3.2 ASIS. ASIS International (in 1955, founded as the nonprofit *American Society for Industrial Security*) certifications are accredited by ANSI.[43]

- *Certified Protection Professional:* The CPP from *ASIS International*[44] (originally called the *American Association for Industrial Security*) demonstrates a thorough understanding of physical, human, and information-security principles and practices. The most senior and prestigious security professional certification covered in this article, the CPP requires extensive on-the-job experience (nine years or seven years with a college degree) including "at least three of those years in responsible charge of a security function," as well as a profound knowledge of technical and procedural security topics and technologies—not just IT security. Only those who have worked with and around security for much of their careers qualify for this credential.[45]
- *Professional Certified Investigator:* The PCI credential focuses on knowledge of legal and evidentiary matters required to present investigations in a court of law, including case management, evidence collection, and case presentation. It requires five years of investigation experience, with at least two years in case management, a high school diploma (or GED), and a clean criminal record.[46]

74.3.3 Brainbench Basic Security Certifications. *Brainbench,* a for-profit private company, offers 640 exams (as of August 2013) in a wide range of subjects including several basic-level security certifications, each requiring the candidate to pass one exam costing about $50.[47] Some of these exams may be useful for new entrants into the job market, for nonprofessionals looking to improve their career prospects,

and for organizations seeking to measure the competence of applicants and employees claiming specific areas of knowledge. A career-development professional wrote about these low-level certifications in 2012,

> You can take a number of exams to meet their requirements on getting certified for a job role. They run specials and offer exams to take for free for a limited time. Most people wait to see when an exam they want to take becomes available for free and then they take the exam. Brainbench offers a yearly membership that cost about $200 dollars [sic] and you can take as many exams as you want. This is good value for the cost and the number of certifications that you can earn.[48]

Examples of these certifications relevant for IA include:

- *Computer Forensics (U.S.):* The Computer Forensics (U.S.) certification is designed for experienced individuals who can prepare for an investigation, analyze and collect evidence, recognize data types, follow proper examination procedures, and initial analysis, use forensic tools, provide initial analysis, and prepare for an investigation and report findings.[49]
- *Firewall Administration Concepts:* "… Audits, Core Technologies, Internals, Maintenance, Network Address Translation (NAT), Optimization, Piercing, Policies, Security Vulnerabilities, and Types."[50]
- *HIPAA (Privacy):* Tests "… knowledge of HIPAA Standards for Privacy. Designed for all professionals who are affected by the HIPAA Standards for Privacy, this test covers the following topics: General Knowledge, Impacted Areas, Implementation, Interactions with External Parties, Regulatory Environment, and Workplace Privacy."[51]
- *HIPAA (Security):* Tests "knowledge of compliance with the standards required by the Security and Electronic Signature Standards as they apply to HIPAA. Designed for healthcare professionals and focusing on nontechnical as well as technical aspects of the HIPAA Standards for Security and Electronic Signatures, this test includes the following topics: Computer Mechanisms, Computer Models, General Knowledge, Implementation, Medical Records, Organization, Security Basics, and Setup."[52]
- *Information Security Administrator:* This certification tests knowledge of networking and Internet security, including authorization, authentication, firewalls, encryption, disaster recovery, and more. Candidates must pass eight exams to obtain this certification.
- *Information Technology Security Fundamentals:* Topics evaluated include "Application Security, Architecture Security, Common Security Services, Data Security, Infrastructure Security, Network Security, Operating System Security, and Threats and Countermeasures."[53]
- *Internet Security:* "Communications, Electronic Mail, Encryption, Identity and Access Management, Malicious Programs, Network Devices and Services, Security Management, Threats, and World Wide Web."[54]
- *ITAA Information Security Awareness:* Information Technology Association of America test for general computer users. Topics include "Computer Best Practices, Computer Ethics & Misuse, ID & Data Information Theft, Internet Best Practices, Passwords, Physical Security, Sensitive Information, and Viruses & Other Harmful Software."[55]

- *Microsoft Security:* Tests "… knowledge of various aspects of security including networks, operating systems, security control concepts, servers, and client computers. Designed for experienced administrators, this test includes the following topics: Networks, Operating Systems, Patch Management, Security Best Practices, Security Controls, Security Incidents, Security Tools, and Servers."[56]
- *Network Authentication:* Tests "… knowledge of authentication protocols, policies, and configuration and audit techniques. Designed for experienced network administrators, this test includes the following topics: Authentication Methods, Encryption, Network Design and Architecture, Protocols, Risk Mitigation, User Accounts, Virtual Private Networking (VPN), and Web Server Authentication."[57]
- *Network Monitoring:* Tests "… knowledge of the core aspects of monitoring and troubleshooting networks. Designed for experienced network administrators, this test covers the following topics: Architecture and Terminology; SNMP Background, Components and Commands; RMON Background, Components and Commands; Other Monitoring and Troubleshooting Tools; System Design and Implementation; and Using SNMP and RMON Tools."[58]
- *Network Security:* Tests "… knowledge of the technical issues involved in maintaining network security. Designed for experienced network administrators, this test covers the following topics: Access Control, Cryptography, Infrastructure, Security Policy and Procedures, System Management, Threats, and Tools."[59]

74.3.4 CyberSecurity Institute. The *CyberSecurity Institute* is a "Noncommercial Research Organization"[60] that provides several courses of interest in addition to certifications:

- Computer Forensics Core Competencies (5 days)[61]
- Computer Forensics for Attorneys (1 day)[62]
- Packet Analysis and Intrusion Detection (5 days)[63]
- Computer Forensics Fundamentals (1 day)[64]
- Windows Forensics Essentials (2 days)[65]

The *CyberSecurity Forensic Analyst* (CSFA)[66] certification aims to identify individuals who can perform a comprehensive and sound forensic examination of a computer system and other digital/electronic devices within a limited time frame. Suggested prerequisites include attendance of the CyberSecurity Institute's *Computer Forensics Core Competencies* course and/or at least one of the following certifications:

- AccessData Certified Examiner (ACE)[67]
- Certified Forensic Computer Examiner (CFCE)[68]
- Certified Computer Examiner (CCE)[69]
- Computer Hacking Forensic Investigator (CHFI)[70]
- EnCase Certified Examiner (EnCE)[71]
- GIAC Certified Forensics Analyst (GCFA)[72]

In addition, candidates should have at least two years of experience performing forensic analysis of Windows FAT and NTFS file systems and writing forensic analysis

reports. Candidates must have no criminal record and must pass an FBI background check.[73]

74.3.5 CERT. *CERT-Certified Computer Security Incident Handler:* The CERT-CCSIH credential is managed by the nonprofit Software Engineering Institute[74] of Carnegie Mellon University, home of the *Computer Emergency Response Team Coordination Center* (CERT/CC).[75]

The CERT-Certified Computer Security Incident Handler (CSIH) certification program has been created for:

- Computer network incident handling and incident responder professionals
- Computer security incident response team (CSIRT) members and technical staff
- System and network administrators with incident handling experience
- Incident handling educators
- Cybersecurity technical staff

The certification is recommended for those computer security professionals with one or more years of experience in incident handling and/or equivalent security-related experience. Although completion of training is not a requirement, we highly recommend that each applicant ensure they have studied appropriately for the examination.[76]

74.3.6 CompTIA. CompTIA is "a nonprofit trade association advancing the global interests of IT professionals and companies …" They "focus our programs on four main areas: education, certification, advocacy, and philanthropy."[77] They provide "four IT certification series that test different knowledge standards, from entry-level to expert."[78]

- Certified Advanced Security Practitioner Certification: CompTIA describes CASP as follows:

  > The CASP certification is an international, vendor-neutral exam that proves competency in enterprise security; risk management; research and analysis; and integration of computing, communications, and business disciplines.
  >
  > The exam covers the technical knowledge and skills required to conceptualize, design, and engineer secure solutions across complex enterprise environments. It involves applying critical thinking and judgment across a broad spectrum of security disciplines to propose and implement solutions that map to enterprise drivers. For more detailed information, download the exam objectives.[79]

- *Cloud+:* This certification "validates the knowledge and best practices required of IT practitioners working in cloud computing environments, who must understand and deliver cloud infrastructure. Recommended experience includes at least 24–36 months of work experience in IT networking, storage, or data center administration, and familiarity with any major hypervisor technologies for server virtualization."[80]
- *Mobile App Security+:* "The idea behind CompTIA Mobile App Security+ is to ensure that developers have the knowledge and skills to design and build applications with security built in from the start." In addition, "There are two editions of CompTIA Mobile App Security+ —one for Android and one for iOS.

Only one exam is required for certification. The exams cover web service and network security, data security and implementing encryption, and other aspects of mobile application security." At the time of writing (August 2013), this exam was under development.[81]

- *Security+:* This security certification focuses on important security fundamentals related to security concepts and theory, as well as best operational practices. Topics include:
  - Access control and identity management
  - Application, data, and host security
  - Compliance and operational security
  - Cryptography
  - Network security
  - Threats and vulnerabilities[82]
- *Social Media Security Professional:* "The SMSP certification designates professionals with demonstrated knowledge of the technical composition of social networking platforms and skills to effectively mitigate risks in order to safeguard organizations' critical information from social media hackers … The SMSP exam is delivered by Ultimate Knowledge Institute (UKI)."[83]

CASP and Security+ are ANSI-accredited.[84]

74.3.7 EC-Council. The for-profit *International Council of E-Commerce Consultants* (EC-Council) was founded in 2001:

> …[EC-Council] is a member-based organization that certifies individuals in various information-security and e-business skills. EC-Council has been certified by American National Standards Institute to meet its ANSI 17024 standard. It is the owner and creator of the world famous Certified Ethical Hacker (CEH), Computer Hacking Forensics Investigator (CHFI) and EC-Council Certified Security Analyst (ECSA)/License Penetration Tester (LPT) programs, and as well as many others programs, that are offered in over 87 countries through a training network of more than 450 training partners globally.
>
> As of January 1st, 2012 EC-Council has trained over 120,000 individuals and certified more than 50,000 security professionals….[85]

Among the certifications managed by EC-Council[86] are the following. Only the C|EH is accredited by ANSI.[87]

- Certified Chief Information Security Officer (C|CISO)
- Certified Disaster Recovery Professional (EDRP)
- Certified Encryption Specialist (E|CES)
- Certified Ethical Hacker (C|EH)
- Certified Hacking Forensic Investigation (C|HFI)
- Certified Incident Handler (E|CIH)
- Certified Network Defense Architect (C|NDA)
- Certified Secure Programmer .NET (E|CSP .NET)
- Certified Security Analyst (E|CSA)
- Certified Security Specialist (E|CSS)

- Certified VoIP Professional (E|CVP)
- Computer Hacking Forensic Investigator (C|HFI)
- Licensed Penetration Tester (L|PT)
- Network Security Administrator (E|NSA)
- Security|5
- Secure Computer User (C|SCU)

74.3.8 CWNP. The *Certified Wireless Security Professional* (CWSP) certification from *Certified Wireless Network Professional* (CWNP),[88] which was founded in 1999, recognizes individuals who can design, implement, and manage wireless LAN security. To obtain this credential, candidates must pass two exams: one for network administration and the other for wireless LAN security.[89]

Other certifications available from CWNP as shown on the cited page include:

- Entry level—CWTS: Certified Wireless Technology Specialist
- Administrator level—CWNA: Certified Wireless Network Administrator
- Professional level—CWDP: Certified Wireless Design Professional
- Professional level—CWAP: Certified Wireless Analysis Professional
- Expert level—CWNE: Certified Wireless Network Expert

74.3.9 GIAC. The for-profit SANS Institute well known for its contributions to research and education in IA.

> The *Global Information Assurance Certification* (GIAC) was founded in 1999 to validate the real-world skills of IT security professionals. GIAC's purpose is to provide assurance that a certified individual has practical awareness, knowledge, and skills in key areas of computer and network and software security. GIAC currently offers certifications for over 20 job-specific responsibilities that reflect the current practice of information security.[90]

A GIAC designation is not a set standard of practiced and learned protocols. Instead SANS offers a number of certifications and certificates that are dedicated to different areas that require information security and different skill sets. Divided into divisions—Audit, Legal, Management, Operations, Security Administration—the certificates are further divided into levels (basic to more advanced skill sets).[91] To attain a GIAC certification, a five- or six-day SANS Training course must be taken followed by an exam. The exam can be taken within a six-month period from the end of the course.

For *Silver Certification*, the exam is taken online and consists of 150 to 200 multiple-choice questions. The cost of a GIAC certification is the same whether the candidate chooses self-study or conference training. Upon successful completion of the exam, a Silver Certification is given to the candidate.[92]

The Silver certification can be further upgraded to a *Gold Certification* two years or more after completing the Silver certification process. Upgrading requires completion of a technical paper related to the candidate's Silver certification. The paper must be completed in a six-month period, although a three-month extension can be added for an extra fee. When the paper is submitted, it is reviewed using a number of criteria:

- Technical—errors
- Education—clarity of thought and advanced knowledge of subject matter

- Extension—a topic outside the box of normal course studies
- Organization—paper clarity, setup, and professional work[93]

The *Global Information Assurance Certification Program* is described by Ed Tittel and Mary Lemons as follows:

> This program seeks to identify individuals who can demonstrate both knowledge of and the ability to manage and protect important information systems and networks. The SANS organization is well known for timely, focused, and useful security information and certification programs. A beacon on this landscape, GIAC now offers regular online classes and uses such classes to draw attendees to their frequent well-situated week-long conferences. Overall, the GIAC program aims at serious, full-time security professionals responsible for designing, implementing, and maintaining a state-of-the-art security infrastructure, which may include incident handling and emergency response team management. Available intermediate-level GIAC credentials include the following:
>
> - GIAC Security Essentials Certification (GSEC)
> - GIAC Information Security Professional (GISP)
> - GIAC Certified ISO-27000 Specialist (G2700)

Available advanced certifications from GIAC include the following 25 (those marked with an asterisk are ANSI-accredited),[94] all of which have links for details on the GIAC home page[91]:

- Assessing and Auditing Wireless Networks (GAWN)
- Certified Enterprise Defender (GCED*)
- Certified Firewall Analyst (GCFW)
- Certified Forensic Analyst (GCFA*)
- Certified Forensic Examiner (GCFE)
- Certified Incident Handler (GCIH*)
- Certified Intrusion Analyst (GCIA*)
- Certified ISO-27000 Specialist (G2700)
- Certified Project Manager (GCPM)
- Certified UNIX Security Administrator (GCUX)
- Certified Web Application Defender (GWEB)
- Certified Windows Security Administrator (GCWN)
- Exploit Researcher and Advanced Penetration Tester (GXPN)
- Information Security Fundamentals (GISF)
- Information Security Professional (GISP)
- Legal Issues in Information Technology & Security (GLEG)
- Mobile Device Security Analyst (GMOB)
- Penetration Tester (GPEN*)
- Reverse Engineering Malware (GREM)
- Secure Software Programmer- .NET (GSSP-.NET)
- Systems and Network Auditor (GSNA*)
- Secure Software Programmer-Java (GSSP-JAVA)

- Security Essentials (GSEC*)
- Security Leadership (GSLC*)
- Web Application Penetration Tester (GWAPT)

A particularly difficult and respected certification is the *GIAC Security Expert* (GSE), which the organization describes as follows:

> The GSE certification is the most prestigious in the IT Security industry. The current exam was developed by subject matter experts and top industry practitioners. The GSE's performance based, hands-on nature sets it apart from any other certifications in the IT security industry. The GSE will determine if a candidate has truly mastered the wide variety of skills required by top security consultants and individual practitioners.
>
> The GSE exam has two parts. The first is a multiple choice exam which may be taken at a proctored location just like any other GIAC exam. The current version of the GSE multiple choice exam has the passing score set at 75% and a time limit of 3 hours. Passing this exam qualifies a person to sit for the GSE hands-on lab. The first day of the two day GSE lab consists of an incident response scenario that requires the candidate to analyze data and report their results in a written report. The second consists of a rigorous battery of hands-on exercises… .[95]

74.3.10 HTCN. The High-Tech Crime Network (HTCN) includes:

> …law enforcement agencies and corporate security professionals from 15 (fifteen) countries. Our membership includes prestigious agencies such as Royal Canadian Mounted Police, Federal Bureau of Investigation, Internal Revenue Service, US Customs, Alcohol Tobacco & Firearms, US Marshals, Department of Defense, National Security Agency, Naval Criminal Investigative Service, Federal Aviation Administration, Department of the Army, and many other law enforcement agencies from Belgium, Canada, England, Germany, Switzerland, and the United States, as well as major telecommunication corporations from across the nation.[96]

HTCN offers four computer-forensic certifications aimed at law enforcement and private-sector IT professionals seeking to specialize in the investigative side of the field.

- Certified Computer Crime Investigator (Basic)
- Certified Computer Crime Investigator (Advanced)
- Certified Computer Forensic Technician (Advanced)
- Certified Computer Forensic Technician (Basic)

The certifications require from three to five years of experience plus 40 to 80 hours of approved training courses plus written reports on from 10 to 40 specific criminal investigations and forensic analyses. Details are available on the organization's Website.[97]

74.3.11 IACIS. The *International Association of Computer Investigative Specialists* (IACIS) offers the *Computer Forensic Computer Examiner* (CFCE) credential to law enforcement and private industry personnel alike. Candidates must have broad knowledge, training, or experience in computer forensics, including forensic procedures and standards, as well as ethical, legal, and privacy issues. Certification requires an intensive peer review, hands-on performance-based testing, as well as a written exam.[68]

A feature of this certification program is its *peer review phase* which must be completed successfully before the candidates can attempt the *Certification phase*:

> The CFCE Program is a two-part process, consisting of a "Peer Review" phase and a "Certification" phase… . Upon entry, each candidate is assigned a "coach" who mentors the candidate through a series of practical exercises (also known as problems) which are based on the CFCE core competencies. Candidates are required to successfully complete the entire peer review phase in order to qualify for the CFCE Certification phase. Candidates who fail the peer review process, or otherwise fail to meet deadlines and other requirements, will not qualify for entry into the CFCE Certification phase.
>
> …The CFCE Certification phase is an independent exercise and not led by a coach. During this phase the candidate is required to demonstrate their knowledge of the CFCE core competencies and practical skills by successfully completing the practical and written examination instruments. Once entry has been approved, the candidate is provided with a forensic image of some type of electronic media (typically a hard disk device) and a practical written examination based on the image (or other media). The candidate must successfully pass the practical examination to qualify for the final examination. Once the candidate has successfully completed the practical examination, they are provided with the final examination. The final examination is a comprehensive written examination consisting of one hundred (100) questions. The questions are based on the CFCE core competencies and deal with a variety of fundamental computer forensic/digital evidence issues.[98]

74.3.12 IIA. The *Institute of Internal Auditors* (IIA) membership (180,000 members worldwide as of August 2013) includes experts "in internal auditing, risk management, governance, internal control, information technology audit, education, and security."

Among the five certifications offered by the IIA is the *Certification in Risk Management Assurance* (CRMA), which began testing in July 2013. The IIA describes the program as follows:

> The CRMA is designed for internal auditors and risk-management professionals with responsibility for and experience in providing risk assurance, governance processes, quality assurance, or control self-assessment (CSA). It demonstrates an individual's ability to evaluate the dynamic components that comprise an organization's governance and enterprise risk-management program and provide advice and assurance around these issues.
>
> The CRMA is one more mark of professional distinction for internal audit practitioners. Earning the CRMA will assist you in demonstrating your ability to:
>
> - Provide assurance on core business processes in risk management and governance.
> - Educate management and the audit committee on risk and risk-management concepts.
> - Focus on strategic organizational risks.
> - Add value for your organization.[99]

74.3.13 ISACA. The nonprofit *ISACA* (formerly known as the *Information Systems Audit and Control Association*)[100] manages several respected certifications. Those marked with an asterisk are ANSI-accredited.

- *Certified in the Governance of Enterprise:* The CGEIT* reflects successful candidates' "knowledge and application of enterprise IT governance principles and practices."[101] ISACA adds that recent studies "ranked CGEIT among the most sought-after and highest-paying IT certifications."[102]
- *Certified Information Systems Auditor:* "The CISA* designation is a globally recognized certification for IS audit control, assurance, and security professionals.

Being CISA-certified showcases your audit experience, skills and knowledge, and demonstrates you are capable to manage vulnerabilities, ensure compliance, and institute controls within the enterprise."[103]

- *Certified Information Security Manager:* "The … CISM* certification promotes international security practices and recognizes the individual who manages designs, and oversees and assesses an enterprise's information security."[104]
- *Certified in Risk and Information Systems Control* (CRISC) was awarded the Best Professional Certification Award from *SC Magazine* in February 2013.[105] "CRISC … is the only certification that positions IT professionals for future career growth by linking IT risk management to enterprise risk management, and positioning them to become strategic partners to the business."[101]

74.3.14 (ISC)^2. The nonprofit *International Information Systems Security Certification Consortium* (ISC)^2 is a nonprofit professional organization[106] providing "vendor-neutral education products, career services, and Gold Standard credentials to professionals in more than 135 countries." As of August 2013, there were 90,000 professionals worldwide certified by (ISC)^2. The organization offers a range of certifications that are accredited under ANSI/ISO/IEC Standard 17024.[107]

The certifications are based on elements of the *Common Body of Knowledge* (CBK) defined by (ISC)^2. Each certification uses a subset of the CBK, which is kept updated yearly by expert committees to adapt to changing environments. The CBK domains are

- Access Control Systems and Methodology
- Application Development Security
- Business Continuity and Disaster Recovery Planning
- Certification and Accreditation (C&A)
- Communications and Network Security
- Cryptography
- Information Security Governance and Risk Management
- Initiate the Preparation Phase
- Law, Investigation, Forensics, and Ethics
- Legal, Regulations, Investigations, and Compliance
- Malicious Code and Activity
- Monitoring and Analysis
- Networks and Communications
- Operations Security
- Perform Execution Phase
- Perform Maintenance Phase
- Physical (Environmental) Security
- Recovery Planning (DRP)
- Risk, Response and Recovery
- Secure Software Concepts
- Secure Software Design

- Secure Software Implementation/Coding
- Secure Software Requirements
- Secure Software Testing
- Security Architecture Analysis
- Security Architecture and Design
- Security Compliance Management
- Security Management Practices
- Security Operations and Administration
- Software Acceptance
- Software Deployment, Operations, Maintenance, and Disposal
- Systems Development Security
- Systems Security Engineering
- Technical Management
- Technology Related Business Continuity Planning (BCP) and Disaster
- Telecommunications and Network Security
- U.S. Government Information Assurance (IA) Governance
- Understand Business Continuity Planning (BCP) and Disaster Recovery Planning (DRP)
- Understand the Purpose of Security Authorization[108]

The certifications offered by $(ISC)^2$ are as follows and generally require a minimum of three years of field experience in information-security–related jobs. The certifications marked by the asterisk (*) are ANSI-accredited.[109]

- Certified Authorization Professional (CAP*)[110]
- Certified Cyber Forensic Professional (CCFPSM)[111]
- Certified Information Systems Security Professional (CISSP*)[112]
- Certified Secure Software Lifecycle Professional (CSSLP*)[113]
- HealthCare Information Security and Privacy Practitioner Information Security and Privacy Practitioner (HCISPP): available November 2013[114]
- Systems Security Certified Practitioner (SSCP*): requires only one year of professional experience in a security-related job[115]

In addition, there are three additional certifications for holders of the CISSP:

- Information Systems Security Architecture Professional (ISSAP*)[116]
- Information Systems Security Engineering Professional (ISSEP*)[117]
- Information Systems Security Management Professional (ISSMP)*[118]

All $(ISC)^2$ certifications require Continuing Professional Education (CPE) credits:

> Credential holders must earn the minimum number of Continuing Professional Education credits (CPEs) annually during each year of the three-year certification cycle. Although members may earn more than the minimum number of CPE credits required for credential maintenance

for the three-year cycle, they are still required to earn and submit the minimum annual number to maintain their certification in "good standing."[119]

CPEs fall into several categories and each certification has different requirements.[120]

Finally, $(ISC)^2$ offers a program for students and other newcomers to the IA field to benefit from membership in and resources of the association. Such candidates can be certified as *Associate of* $(ISC)^2$ for the CAP, CCFP, CISSP, CSSLP, or SSCP certifications by passing the appropriate exam, paying a modest fee, and accumulating the appropriate number of CPE credits. After the requisite number of years of experience have been accumulated, the *Associate* status is converted to full certification.[121]

74.3.15 mile2. The for-profit mile2 training firm offers a "…worldwide network of Authorized Training Partners provides Cybersecurity training and consulting services that meet military, government and corporate specifications." Certifications listed on their home page include:

- C)DFE Certified Digital Forensics Examiner
- C)DRE Certified Disaster Recovery Engineer
- C)IHE Certified Incident Handling Engineer
- C)ISSO Certified Information Systems Security Officer
- C)NFE Certified Network Forensics Examiner
- C)PTC Certified Penetration Testing Consultant
- C)PTE Certified Penetration Testing Engineer
- C)SLE Certified Secure Linux Engineer
- C)SLO Certified Security Leadership Officer
- C)SS Certified Security Sentinel
- C)SWAE Certified Secure Web Applications Engineer
- C)VA Certified Vulnerability Assessor
- C)WSE Certified Wireless Security Engineer
- IS20 Security Controls
- ISCAP Information Systems Certification and Accreditation Professional[122]

74.3.16 Security University. The for-profit *Security University* (SU) was established in 1999 and offers several security certifications and a wide range of classes, most of which involve hands-on training.

Security University is accredited by ACCET, the *Accrediting Council for Continuing Education and Training*; ACCET is listed by the U.S. Department of Education as a nationally recognized accrediting agency and is a recognized member of the *Council for Higher Education Accreditation*. Security University is an institute of higher education in the State of Virginia certified to operate by the *State Council of Higher Education for Virginia*. SU classes and certifications are GI Bill Approved. The National Security Agency has authorized Security University to award the *CNSS National Standard 4011, 4012, 4013, 4015*, and *4016E, I* and *A* certifications.

- *Q/ISP Qualified/ Information Security Professional:* complies with CNSS 4011/ 4102/4013/4015/4016A standards. Requires demonstration of both knowledge

and proficiency in information security, including hands-on skills. Maintenance of the certification requires "participation in research or study, attendance at recognized subject-matter training and professional educational programs, presentation or publication of information security papers, contributions to the information security Q/ISP Body of Knowledge, or service in professional organizations."[123]

- *Q/IAP Qualified/ Information Assurance Professional:* for IT security professionals, system administrators, auditors, CISOs, and others "looking to build tactical security skills and improve their career and income."[124] This certification requires successful completion of three SU courses:
  - Qualified Access, Authentication & PKI Professional (a hands-on course)
  - Qualified Security Policy Admin & (SOA) Security Oriented Architect
  - Qualified Security Certification and Accreditation Process
- *Q/SSE Qualified/Software Security Expert:* This certificate is for everyone concerned with securing program code[125]
- *Q/WNP™ Wireless Certifications:* "The CWTS certification is the entry-level wireless certification for the IT industry designed for WLAN sales and technical support professionals. This certification will get you started in your IT career by ensuring you have a solid base of applicable knowledge of: fundamentals of RF behavior, can describe the features and functions of wireless components, and have the skills needed to install and configure wireless network hardware components."[126]

Security University has a complete list of its 60 courses (most of them 40 hours or more) available in its online catalog, with details available for every course.[127]

74.4 PREPARING FOR SECURITY-CERTIFICATION EXAMS. The key to passing the security-certification exams is daily attention to expanding one's exposure to interesting and thought-provoking information and ideas in the field. Cramming does not work; it is not possible to remember what is learned in a rush for very long. Indeed, all students should learn to use SQ3R (survey/question, read/recite, review) a well-established study method that pays off with long-term integration and retention of knowledge.

Even when participating in serious, extended courses lasting many days, participants should review constantly and challenge themselves with questions about concepts, terminology, and subtleties. If the course material or the instructor's comments don't make sense on a particular point, the student must not let the mystery pass, but should question authority if necessary to ensure that every point is understood and integrated into a coherent picture of the issues. Students can also look online for recorded lectures[128] and slide presentations from academic or professional meetings and courses.[129]

Granted, some test takers may enroll in short courses (boot camps) devoted to preparing for specific certifications, but without long-term commitment to professionalism, such efforts result in a certificate that raises eyebrows among employers and clients when the crammer fails to display an ingrained understanding of soon-forgotten concepts and terminology.

Anyone committed to professionalism and long-term success in a field should read a wide range of reputable publications and participate in serious discussion groups where participants exchange real-world experiences and analysis.

74.5 CONCLUDING REMARKS. This chapter cannot provide an exhaustive listing of training and certification resources. Practitioners should monitor the trade press for indications of trends in acceptance of new certifications and should judge how to keep themselves at the forefront of the field through continued professional development. In addition, professionals may find it satisfying, personally educational, and even moderately lucrative to involve themselves in training and certification as teachers and proctors.

74.6 FURTHER READING

@RISK: The Consensus Security Vulnerability Alert: www.sans.org/newsletters/risk

Armstrong, I., ed. *SC Magazine.* www.scmagazine.com/#

Bunting, S. *EnCase Computer Forensics—The Official EnCE: EnCase Certified Examiner Study Guide*, 3rd ed. Sybex, 2012.

Chapple, M. *The GSEC Prep Guide: Mastering SANS GIAC Security Essentials*. Wiley, 2003.

Computer Security Portal from the National Institute of Standards and Technology Information Technology Laboratory Computer Security Division's Computer Security Resource Center: www.nist.gov/computer-security-portal.cfm

Computerworld Newsletters: www.computerworld.com/s/newsletters

Conrad, E., S. Misenar, and J. Feldman. *CISSP Study Guide*, 2nd ed. Syngress, 2012.

Conrad, E., S. Misenar, and J. Feldman. *Eleventh Hour CISSP Study Guide*, 2nd ed. Syngress, 2013.

CRYPTO-GRAM from Bruce Schneier: www.schneier.com/crypto-gram.html

Daily Open Source Infrastructure Report: https://www.dhs.gov/dhs-daily-open-source-infrastructure-report

Docter, Q., E. Dulaney, and T. Skandier. *CompTIA A+ Complete Deluxe Study Guide Recommended Courseware: Exams 220-801 and 220-802*, 2nd ed. Sybex, 2012.

Docter, Q., E. Dulaney, and T. Skandier. *CompTIA A+ Complete Study Guide Authorized Courseware: Exams 220-801 and 220-802* 2nd ed. Sybex, 2012.

EFFector from the Electronic Frontier Foundation: www.eff.org/effector

EPIC Alert from the Electronic Privacy Information Center: www.epic.org/alert

Harris, S. *CISSP All-in-One Exam Guide*, 6th ed. McGraw-Hill Osborne Media, 2012.

Harris, S. *CISSP Boxed Set*, 2nd ed. (All-in-One). McGraw-Hill Osborne Media, 2012.

Hernandez, S., ed. *Official (ISC)^2 Guide to the CISSP CBK,* 3rd ed. Auerbach, 2012.

Information Systems Security Association. *The ISSA Journal.* www.issa.org/?page=ISSAJournal

ITL Security Bulletins from the National Institute of Standards and Technology Information Technology Laboratory Computer Security Division's Computer Security Resource Center http://csrc.nist.gov/publications/PubsITLSB.html

Krutz, R. L., and A. J. Fry. *The CSSLP Prep Guide: Mastering the Certified Secure Software Lifecycle Professional.* Wiley, 2009.

Lammle, T. *CompTIA Network+ Study Guide Authorized Courseware: Exam N10-005*, 2nd ed. Sybex, 2012.

Miller, L., and P. H. Gregory. *CISSP for Dummies*, 4th ed. For Dummies, 2012.

Network World Security Research Center: www.networkworld.com/topics/security.html

RISKS FORUM DIGEST from the Association for Computing Machinery Committee on Computers and Public Policy: http://catless.ncl.ac.uk/Risks

SANS newsletters: www.sans.org/newsletters

Special Publications (800 Series) from the National Institute of Standards and Technology Information Technology Laboratory Computer Security Division's Computer Security Resource Center: http://csrc.nist.gov/publications/PubsSPs.html

Stewart, J. M., M. Chapple, and D. Gibson. *CISSP: Certified Information Systems Security Professional Study Guide*, 6th ed. Sybex, 2012.

Walker, M. *CEH Certified Ethical Hacker Boxed Set* (All-in-One). McGraw-Hill Osborne Media, 2013.

74.7 NOTES

1. DHS Press Release, "Secretary Napolitano Announces New Hiring Authority for Cybersecurity Experts," October 1, 2009, www.dhs.gov/news/2009/10/01/secretary-napolitano-announces-new-hiring-authority-cybersecurity-experts
2. National Security Council, "Cybersecurity," May 29, 2009, www.whitehouse.gov/administration/eop/nsc/cybersecurity (URL inactive).
3. E. Engleman, K. Riddel, A. Holmes, and Raytheon, "SAIC Vie with U.S. in 'Fratricide' Over Cyberworkers." *Bloomberg News*, March 29, 2011 (reference #13 in original document), www.bgov.com/news_item/HMRFKLKr4Dj_sBtkxFos-Q (URL inactive) (registration required).
4. Homeland Security Advisory Council, "CyberSkills Task Force Report, Fall 2012," www.dhs.gov/sites/default/files/publications/HSAC%20CyberSkills%20Report%20-%20Final.pdf
5. B. Ballenstedt, "Eleven Tips for Building a Stronger Cybersecurity Workforce," *Nextgov Wired Workplace*, October 3, 2012, www.nextgov.com/cio-briefing/wired-workplace/2012/10/eleven-tips-building-stronger-cybersecurity-workforce/58564/
6. *Microsoft Encarta 2008*. © 1993–2007 Microsoft Corporation.
7. J. Walsh, "6 Essential Skills for Cyber Security Professionals," AE ashley ellis, August 2, 2013, www.ashleyellis.com/2013/08/6-essential-skills-for-cybersecurity-professionals
8. A one-page summary of the useful Survey-Question-Read-Recite-Review (SQ3R) technique is available from one of the authors: www.mekabay.com/methodology/sq3r.pdf
9. A behaviorally defined learning objective is something that can be observed or measured; for example, "understanding firewall configuration parameters" is not a behavioral objective, whereas "configuring firewalls to meet defined goals" is a behavioral objective.
10. I. Armstrong, "Money Matters: SC Magazine/EC-Council Salary Survey 2007," *SC Magazine*, April 1, 2007, www.scmagazineus.com/Money-matters-SC-MagazineEC-Council-Salary-Survey-2007/article/34874/
11. M. Savage, "Get Qualified: Certification—That's the Name of the Game," *SC Magazine*, November 1, 2003, http://scmagazine.com/asia/news/article/419631/get-qualified-certification-thats-name-game (URL inactive).
12. A. Saita, "The CISSP Receives International Standardization," *SearchSecurity.com*, June 23, 2004, http://searchsecurity.techtarget.com/original-Content/0,289142,sid14_gci989901,00.html (URL inactive).
13. $(ISC)^2$, "DoD Fact Sheet," 2007, https://www.isc2.org/dod-fact-sheet.aspx?terms=DoD%20Fact%20Sheet%2c%202007

14. C. Hildebrand, "Forecast 2008: Information Security Goes Strategic," *InfoSecurity Professional* (Spring 2008): 4–7, www.isc2.org/uploadedFiles/(ISC)2_Member_Content/Member_Resources/infosecuritypro2008spring(1).pdf
15. L. G. Paul, "How Valuable Are Security Certifications Today?" *CSO Security Leadership*, April 1, 2003, www.csoonline.com/article/731010/how-valuable-are-security-certifications-today-
16. *2011 (ISC)2 Global Information Security Workforce Study* p 22. https://www.isc2.org/uploadedFiles/Industry_Resources/FS_WP_ISC%20Study_020811_MLW_Web.pdf
17. *2011 (ISC)2 Global Information Security Workforce Study*, p. 23.
18. *2011 (ISC)2 Global Information Security Workforce Study*, Fig. 27.
19. *2013 (ISC)2 Global Information Security Workforce Study* p 14. https://www.isc2.org/uploadedFiles/(ISC)2_Public_Content/2013%20Global%20Information%20Security%20Workforce%20Study%20Feb%202013.pdf
20. M. Suby. Personal communication, August 16, 2013.
21. "Accreditation Program for Personnel Certification Bodies under ANSI/ISO/IEC 17024," ANSI Website, https://www.ansica.org/wwwversion2/outside/PERgeneral.asp; see also E. Gasiorowski-Denis, "New and Improved ISO/IEC 17024 Standard for Personnel Certification Programmes," ISO Website, July 24,2012, www.iso.org/iso/home/news_index/news_archive/news.htm?refid=Ref1625
22. "ISO/IEC 17011:2004 Conformity Assessment—General Requirements for Accreditation Bodies Accrediting Conformity Assessment Bodies," ISO Website, Standard reviewed and confirmed in 2008, www.iso.org/iso/catalogue_detail?csnumber=29332
23. Institute for Credentialing Excellence, www.credentialingexcellence.org
24. NCCA Accreditation, www.credentialingexcellence.org/ncca
25. American Board for Certification in Homeland Security, www.abchs.com
26. American Health Information Management Association, www.ahima.org
27. Defense Security Service, www.cdse.edu/about.html
28. CNSS, "CNSS History," https://www.cnss.gov/history.html
29. CNSS,"Frequently Asked Questions | What are National Security Systems?" https://www.cnss.gov/faq.html
30. CNSS, CNSSI 4009, "National Information Assurance Glossary," CNSS Instruction No. 4009 (2010-04-26), https://www.cnss.gov/Assets/pdf/cnssi_4009.pdf
31. CNSS, "Unclassified Index," https://www.cnss.gov/full-index.html
32. NSTISSI 4011, https://www.cnss.gov/Assets/pdf/nstissi_4011.pdf
33. CNSSI 4012, https://www.cnss.gov/Assets/pdf/cnssi_4012.pdf
34. CNSSI 4013, https://www.cnss.gov/Assets/pdf/cnssi_4013.pdf
35. CNSSI 4014, https://www.cnss.gov/Assets/pdf/cnssi_4014.pdf
36. NSTISSI 4015, https://www.cnss.gov/Assets/pdf/nstissi_4015.pdf
37. CNSSI 4016, https://www.cnss.gov/Assets/pdf/CNSSI-4016.PDF
38. J. Vijayan, "Demand for IT Security Experts Outstrips Supply: Employers Will Pay More for Certified—and Experienced—IT Security Pros, Studies Find," *Computerworld*, March 7, 2013, www.computerworld.com/s/article/9237394/Demand_for_IT_security_experts_outstrips_supply

39. A particularly valuable publication that was immensely helpful in organizing the following information was written by Ed Tittel and Mary Lemons in May 2013. The authors express their gratitude to these colleagues. Reference: E. Tittel and M. Lemons, "SearchSecurity.com's IT Security Certifications Guide," *SearchSecurity*, May 2013, http://searchsecurity.techtarget.com/tip/SearchSecuritycom-guide-to-information-security-certifications
40. ACFE, "About the ACFE," www.acfe.com/about-the-acfe.aspx
41. ACFE, "2013/2014 Compensation Guide for Anti-Fraud Professionals." http://www.acfe.com/uploadedFiles/ACFE_Website/Content/documents/2013-comp-guide.pdf
42. ACFE, "The CFE Credential," www.acfe.com/cfe-credential.aspx
43. ANSI Accreditation Services: "ASIS International," https://www.ansica.org/wwwversion2/outside/ALLdirectoryDetails.asp?menuID=2&prgID=201&orgID=90&status=4
44. ASIS International, "About ASIS," https://www.asisonline.org/About-ASIS/Pages/default.aspx
45. ASIS International, "Certified Protection Professional," https://www.asisonline.org/Certification/Board-Certifications/CPP/Pages/default.aspx
46. ASIS International, "Professional Certified Investigator," https://www.asisonline.org/Certification/Board-Certifications/PCI/Pages/default.aspx
47. Brainbench Test Center, www.brainbench.com/xml/bb/common/testcenter/alltests.xml
48. G. Singh, "Boost Your Resume with Brainbench Certifications," *Examiner.com*, September 20, 2012, www.examiner.com/article/boost-your-resume-with-brainbench-certifications
49. Brainbench, "Computer Forensics (U.S.)," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=2516
50. Brainbench, "Firewall Administration Concepts," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=2581
51. Brainbench, "HIPAA (Privacy)," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=1489
52. Brainbench, "HIPAA (Security)," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=1143
53. Brainbench, "Information Technology Security Fundamentals," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=1542
54. Brainbench, "Internet Security," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=110
55. Brainbench, "ITAA Information Security Awareness," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=1349
56. Brainbench, "Microsoft Security," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=2723
57. Brainbench, "Network Authentication," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=2636
58. Brainbench, "Network Monitoring," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=424
59. Brainbench, "Network Security," http://brainbench.com/xml/bb/common/testcenter/consumer/taketest.xml?testId=30

60. CyberSecurity Institute, www.csisite.net
61. CyberSecurity Institute, "Digital Forensic Examiner," www.csisite.net/training/core.htm
62. CyberSecurity Institute, "Computer Forensics for Attorneys," www.csisite.net/training/attorneys.htm
63. CyberSecurity Institute, "Packet Analysis and Intrusion Detection," www.csisite.net/training/packetanalysis.htm
64. CyberSecurity Institute, "Computer Forensics Fundamentals," www.csisite.net/training/computerforensicsfundamentals.htm
65. CyberSecurity Institute, "Windows Forensics Essentials," www.csisite.net/ training/wfe.htm
66. CyberSecurity Institute, "CSFA," www.cybersecurityforensicanalyst.com
67. AccessData, www.accessdata.com/training/certifications
68. International Association of Computer Investigative Specialists (IACIS), https://www.iacis.com/certification
69. International Society of Forensic Computer Examiners (ISFCE), https://www.isfce.com/certification.htm
70. EC-Council, www.eccouncil.org/certification/computer-hacking-forensics-investigator
71. Guidance Software, www.guidancesoftware.com/computer-forensics-training-ence-certification.htm
72. Global Information Assurance Certification (SANS Institute), www.giac.org/certification/certified-forensic-analyst-gcfa
73. CyberSecurity Institue, "FBI Background Check," www.cybersecurityforensicanalyst.com/background.htm
74. Software Engineering Institute, www.sei.cmu.edu
75. CERT® Coordination Center (CERT/CC), www.cert.org/certcc.html
76. Software Engineering Institute, "CERT-Certified Computer Security Incident Handler," www.sei.cmu.edu/certification/security/csih
77. CompTIA, "About CompTIA," www.comptia.org/aboutus.aspx
78. CompTIA, "CompTIA Certifications," http://certification.comptia.org/getCertified/certifications.aspx
79. CompTIA, "CompTIA Advanced Security Practitioner," http://certification.comptia.org/getCertified/certifications/casp.aspx
80. CompTIA, "CompTIA Cloud+," http://certification.comptia.org/getCertified/certifications/cloudplus.aspx
81. CompTIA, "CompTIA Mobile App Security+," http://certification.comptia.org/getCertified/certifications/mobile_app_security.aspx
82. CompTIA, "CompTIA Security+," http://certification.comptia.org/getCertified/certifications/security.aspx
83. CompTIA, "Social Media Security Professional," http://certification.comptia.org/getCertified/certifications/smsp.aspx
84. ANSI Accreditation Services, "CompTIA," https://www.ansica.org/wwwversion2/outside/ALLdirectoryDetails.asp?menuID=2&prgID=201&orgID=93&status=4

85. EC-Council, "About EC-Council," www.eccouncil.org/about-us
86. EC-Council, "Certification," www.eccouncil.org/Certification
87. ANSI Accreditation Services, "International Council of E-Commerce Consultants (EC-Council)," https://www.ansica.org/wwwversion2/outside/ALLdirectoryDetails.asp?menuID=2&prgID=201&orgID=211&status=4
88. CWNP, "About CWNP," www.cwnp.com/about
89. CWNP, "About CWSP," www.cwnp.com/certifications/cwsp
90. SANS, "Certification," www.sans.org/help/site-network
91. GIAC, "Global Information Assurance Certification," www.giac.org
92. Jeff Pike, "GIAC: The Hands-On IT Security Certification," *Certification Magazine*, n.d., www.certmag.com/read.php?start=0&in=3469
93. GIAC, "Certifications: Gold Overview," www.giac.org/certifications/gold
94. ANSI Accreditation Services, "Global Information Assurance Certification (GIAC)," https://www.ansica.org/wwwversion2/outside/ALLdirectoryDetails.asp?menuID=2&prgID=201&orgID=98&status=4
95. GIAC, "GIAC Security Expert (GSE) Overview," www.giac.org/certification/security-expert-gse
96. HTCN, "History," www.htcn.org
97. HTCN, "Certification Requirements," www.htcn.org/cert.htm#
98. IACIS, "Overview of the CFCE Program," https://www.iacis.com/certification/external_overview
99. IIA, "Certification in Risk Management Assurance," https://na.theiia.org/certification/CRMA-Certification/Pages/CRMA-Certification.aspx
100. ISACA, "History of ISACA," www.isaca.org/About-ISACA/History/Pages/default.aspx
101. ISACA, "ISACA Certification: IT Audit, Security, Governance and Risk," www.isaca.org/CERTIFICATION/Pages/default.aspx
102. ISACA, "Foote Partners Report: ISACA Certifications Earn Top Pay Premiums," Press Release, September 8, 2011, www.isaca.org/About-ISACA/Press-room/News-Releases/2011/Pages/Foote-Partners-Report-ISACA-Certifications-Earn-Top-Pay-Premiums.aspx
103. ISACA, "Certified Information Systems Auditor (CISA)," www.isaca.org/Certification/CISA-Certified-Information-Systems-Auditor/Pages/default.aspx
104. ISACA, "Certified Information Security Manager," www.isaca.org/Certification/CISM-Certified-Information-Security-Manager/Pages/default.aspx
105. ISACA, "ISACA's CRISC Certification Wins Best Professional Certification Award from SC Magazine," Press Release, February 27, 2013, www.isaca.org/About-ISACA/Press-room/News-Releases/2013/Pages/ISACA-CRISC-Certification-Wins-Best-Professional-Certification-Award.aspx
106. The acronym is a geek joke: The original acronym was IISSCC which someone collapsed into the algebraic equivalent $(ISC)^2$ shortly after the organization was founded in1989.
107. $(ISC)^2$. "About $(ISC)^2$," https://www.isc2.org/aboutus/default.aspx
108. $(ISC)^2$. "The Compendium of Information Security Topics," https://www.isc2.org/uploadedFiles/Education/Review_Seminars/CBK.pdf

109. ANSI Accreditation Services, "International Information Systems Security Certification Consortium, Inc. $(ISC)^2$," https://www.ansica.org/wwwversion2/outside/ALLdirectoryDetails.asp?menuID=2&prgID=201&orgID=97&status=4
110. $(ISC)^2$, "CAP," https://www.isc2.org/cap/default.aspx
111. $(ISC)^2$, "CCFP," https://www.isc2.org/ccfp/default.aspx
112. $(ISC)^2$, "CISSP," https://www.isc2.org/cissp/default.aspx
113. $(ISC)^2$, "CSSLP," https://www.isc2.org/csslp/default.aspx
114. $(ISC)^2$, "HCISPP," https://www.isc2.org/hcispp/default.aspx
115. $(ISC)^2$, "SSCP," https://www.isc2.org/sscp/default.aspx
116. $(ISC)^2$, "About ISSAP," https://www.isc2.org/issap.aspx
117. $(ISC)^2$, "About ISSEP," https://www.isc2.org/issep.aspx
118. $(ISC)^2$, "About ISSMP," https://www.isc2.org/issmp/default.aspx
119. $(ISC)^2$, "Maintaining Your Credential," https://www.isc2.org/maintaining-your-credential.aspx
120. $(ISC)^2$, "Continuing Professional Education," https://www.isc2.org/uploadedFiles/Credentials_and_Certifcation/About_Our_Credentials_and_Process/CPE.pdf (URL inactive).
121. $(ISC)^2$,"Associate of $(ISC)^2$," https://www.isc2.org/associates/default.aspx
122. mile2 home page, http://mile2.com
123. Security University, "What is a Q/ISP 'Qualified' Information Security Professional?" www.securityuniversity.net/QISP.php
124. Security University, "What is a Q/IAP 'Qualified" Information Assurance Professional?" www.securityuniversity.net/QIAP.php
125. Security University, "Q/SSE Qualified/ Software Security Expert Certification," www.securityuniversity.net/QSSE.php
126. Security University, "CWTS Wireless LAN Technology," www.securityuniversity.net/classes_wireless_cert.php
127. Security University, "Roadmap," www.securityuniversity.net/classes.php
128. A useful tool when searching YouTube for lectures on specific topics is to append ", long" to the search string. For example, searching for "authentication" on YouTube in August 2013 resulted in 87,400 hits, but "authentication, long" resulted in 4,230 results, most of which were 20 minutes or longer.
129. See, for example, Professor Kabay's freely available course materials, www.mekabay.com/courses/index.htm

CHAPTER 75

THE FUTURE OF INFORMATION ASSURANCE

Jeremy A. Hansen

75.1 INTRODUCTION 75·1

75.2 COMPLEXITY 75·1

75.3 CHANGING POSTURE 75·2

75.3.1 Quarantine, DMZ, and Honeycombs 75·3

75.3.2 BYOD 75·3

75.4 SOME SPECIFIC SLICES OF THE FUTURE 75·4

75.4.1 Quantum Computing 75·4

75.4.2 Healthcare 75·5

75.4.3 Legal & Jurisdictional Issues 75·5

75.4.4 Ubiquitous Technologies 75·7

75.5 NEW DIRECTIONS 75·8

75.6 NOTES 75·9

75.1 INTRODUCTION. As previous installments of this final chapter have done in other editions of this *Handbook*, I take a moment to tip my hat to the authors of the previous chapters for their informative coverage of topics in information assurance (IA). This chapter will be much less precise than earlier chapters for reasons that will become obvious, but hopefully no less informative.

In a departure from the typical academic style that I am comfortable with, I hope to dabble a bit in the dangerous game of predicting the future. I'll also make this a highly personal work and will address the reader directly.

Attempts to predict the future tend to extrapolate a spectrum of outcomes that we, residents of the present, should look out for. Others try to be more specific and try to lead the target and guess a single plausible outcome, like a hunter taking a careful shot at a running deer. I intend to find a middle ground here by considering specific possible outcomes, but select them in a spectrum narrow enough to group them together in a loosely logical whole.

75.2 COMPLEXITY. It is important to distinguish here between the real world and the constrained world of information assurance (IA) that appears in many academic sources, popular security sites, and best practices. Like many scientific treatments, the constrained world in its many forms requires us to make certain assumptions for certain conclusions to be true. If the compression function of a cryptographic hash algorithm acts like a random oracle, then we can conclude that some specified construction is secure. This is useful! The trouble with systems with constraints and assumptions is that

they are often pressed into service as subsystems of a larger solution. Even then, they can be useful, but may now be working in applications where their constraints are violated and their assumptions no longer hold. The more complex the systems become, the more buried the subsystems become. At a certain point of high complexity, the number of vulnerabilities of a system becomes greater than the sum of the vulnerabilities of its parts.

We live in a complex civilization, with increasingly complex interconnections between the different moving parts: government, industry, utilities, technological advances, culture, and so on. The cultural, physical, and electronic infrastructure required to maintain what we have is immense.

Here is a thought exercise: look to your left and pick out something you see. Consider its component parts and the process required to construct it. Now pick any one of those component parts or one of the technologies used in its construction. Consider its component parts and the process required to construct it. Do this recursively until you get bored or happen upon some primitive component.

Consider how many of the processes and technologies that you identified in that technological family tree that we no longer have access to. Our modern technologies are being built on an increasingly narrow set of parent technologies, the loss of which could shake the supply chain of the descendant technologies.

Now we circle back to IA. The information systems that we have constructed are brittle and porous, no matter how many layers of steel we bolt on or how many fingers we hire to press in the holes of the dam. It is not possible to secure perfectly all the various component parts of the systems we rely on. We spend an awful lot of money doing this, and spending money protecting our systems with brittle and porous countermeasures will continue to be necessary (but not sufficient) for the foreseeable future.

And that is OK.

It is OK because you are already compromised, or will be shortly. Perhaps not in a big way, perhaps not in a detectable way, it has happened or it will happen. I am asking you for a moment to take a deep breath and let the waves of terror at the thought of this drift away. I am asking you to sit up straighter, take another breath, and change your mental posture when it comes to information assurance. Accept that you have been or will be compromised.

Changing one's (or one's organization's) security posture will likely be awkward and greeted with suspicion. Embracing the possibility (indeed, the likelihood) of failure can be perceived as cowardly. However unpalatable the idea of compromise, not having a post-compromise plan of action is the sort of irresponsibility that can sink an organization. On a personal level, if you have a lot of assets, it makes sense to have a will! This change in posture is also not likely to please developers of security software and devices—silver-bullet security promises move much more product and result in more subscribers per month than the compromise-anticipating approach.

Please do not misunderstand—you should not run out and turn off all your antivirus software to free up the processor cycles. As a matter of fact, your antivirus software is a sensible place to start, since its quarantine offers an easily understood approach to accepting compromise that I'll explore shortly.

75.3 CHANGING POSTURE. An attacker usually must expend fewer resources to compromise a target than a defender must expend defending that target. Realizing this cost asymmetry must force us to reconsider our approach to protecting information. Like previous attempts at handling asymmetric threats, we should not expect our

approach to be successful all of the time. Nor should we expect that simply spending more money will necessarily lead to better results.

Your organization does not and cannot spend enough money on information security. Some managers, directors, and C-level executives realize this and are satisfied with *good enough* (not that they would probably say so), but are often aggressively confronted by employees and vendors with horror stories and dire warnings. You might know this as fear, uncertainty, and doubt, or FUD, and most would agree that though it can be effective, it is inappropriate and can distract from proper allocation of resources. Because resources are limited (and in spite of the cloud of FUD), one basic question needs to be answered to evaluate your current and future security postures: *What is your information security team's job?*

What is it now and what should it be? Is it to prevent all threats, external and internal (or foreign and domestic, if you prefer)? Is it to ensure that PCI, HIPAA, or whatever regulatory compliance checklist applies to you stays completely checked? Is it to avoid a public breach?[1]

75.3.1 Quarantine, DMZ, and Honeycombs. You don't get a towel-wrapped bottle of champagne, party favors, and an oversized check that says "Congratulations! You're compromised!" Almost everyone can agree that a compromise is a demoralizing event, usually resulting in a period of uncertainty and a raft of questions: Why didn't our PCI auditor point this out? Why did our Senior Technical Analyst put that USB drive into his system while he was running the weekly batch job? Where else did the bad actors go on the network?

Answers to the first two questions are not particularly helpful. The last question is often the most troubling and long-lasting. It is also a question you should spend some resources answering with a proper post-mortem. Once you have the answer (as little consolation as it might be), consider what you could have done to (a) make the post-mortem easier and (b) how you could have more effectively restricted the movements of the intruder.

Making the post-mortem easier requires good data-collection and analysis tools, which are available, but may not be accessible to organizations with small budgets. Much more interesting is how intruders can be confined to a small portion of the network, database, or application space without compromising usability by legitimate users. The demilitarized zone (DMZ) model, where Internet-facing servers are segregated from the internal network, is quite popular, but reinforces the *hard-crunchy-shell-with-a-soft-chewy-center* architecture. A better architecture would have pockets of chewy where access to related resources can be shared or simplified (using single sign on or the like), with crunchy boundaries between these pockets. I envision this as a honeycomb, with access compartmentalized inside each cell of the honeycomb, and each cell separated by rigid firewall rules or other access control. Compromising one of the cells, as damaging as that might be, would not lead to an organization-wide crisis. When compromise is detected, the affected cell may be quarantined from the rest of the network until a proper response can be mobilized.

A similar approach is used in companies that connect to the networks of partners via virtual private networks (VPNs). Legitimate bidirectional access is necessary, but restrictions on traffic entering the organization's network keep a partner's compromise from spilling over.

75.3.2 BYOD. Another effect of tightly segregated access is to mitigate the security concerns of a permissive Bring-Your-Own-Device (BYOD) policy. When a

company supports BYOD, its information technology (IT) infrastructure is expected to allow, if not support, personally owned devices on the corporate network. Naturally, it gives security folks a headache. Taking BYOD to a logical (but conceivable) extreme, let's look at a somewhat ridiculous expectation for a network: BED, or Bring Every Device, where a network accepts and offers connectivity to every device requesting such access. Note that this is how many Internet service providers operate as impartial carriers of data.

BED means that devices known to be malicious will use the network. Trusted devices will use the network, too, and they might receive preferential access even before any user authentication takes place. In fact, separating device authentication from user authentication provides for access to legitimate users who may be using an untrusted personal device, and can restrict untrusted users who gain access to a trusted device.

Current network access control (NAC) techniques seem to be part of the solution, as they allow for authentication of devices and can properly isolate unauthenticated devices. The rest of the solution includes integrating NAC more tightly with LDAP or Kerberos in order to construct the requisite cells of the honeycomb architecture.

75.4 SOME SPECIFIC SLICES OF THE FUTURE.

With the *change your security posture* lecture over, I'd like to reflect on a few areas that I believe will become more important or change the way we think about assurance in the future:

- Quantum computing
- Healthcare
- Legal and jurisdictional issues
- Ubiquitous technologies

75.4.1 Quantum Computing.

This field is poorly understood by the majority of IT and security practitioners, and not well understood by most computer scientists. There is a good reason for this: It involves a physics and mathematics (primarily linear algebra) background to manage. Learning how these systems operate, much less how to secure them, will remain a formidable challenge should quantum computers see widespread deployment. I will provide an extremely brief overview here (paraphrased from Scott Aaronson's 2008 *Scientific American* article), but suggest Michael Nielsen's terrific video course[2] for those interested in more depth.

The basic idea of quantum computing is that subatomic particles relate to other particles in interesting and useful ways. If we have 1000 of these particles, they can be in a superposition that can represent 2^{1000} (about 10^{300} if you prefer that notation) different values—simultaneously.

I'll let Scott Aaronson take over:

> Then, by performing various operations on the particles and on some auxiliary ones—perhaps hitting them with a sequence of laser pulses or radio waves—we can carry out an algorithm that transforms all 10^{300} numbers (each one a potential solution) at the same time. If at the end of doing that we could read out the particles' final quantum state accurately, we really would have a magic computer: it would be able to check 10^{300} possible solutions to a problem, and at the end we could quickly discern the right one.[3]

There are catches. Among them, it is only possible to measure one of these solutions and there is no certainty that a computer operating on 1000 particles can even be built. Should these obstacles be overcome, one of the best-known security implications

of such computers is that cryptographic algorithms that rely on factoring integers or calculating the discrete logarithm may be dramatically weakened. *Post-quantum cryptography* is a reasonably new[4] subfield of cryptography that investigates cryptographic systems that can remain strong even when attackers have quantum computers.

75.4.2 Healthcare. Consider healthcare security as broadly having two components:

1. Medical records and medical data security
2. Medical systems security

Though the Health Information Technology for Economic and Clinical Health (HITECH) Act and Health Insurance Portability and Accountability Act of 1996 (HIPAA) are not new or remotely futuristic, the implications of moving to an all-digital medical records system are. Medical records are inextricably linked to a single patient, so a disclosure of any sort can linger with that patient for the rest of his or her life. Chapter 71 in this *Handbook*, "Medical Records Security" covers this topic in much more detail.

The security of medical systems, on the other hand, is among the most interesting (to me) topics in the future of security. As computers become smaller and medical devices more sophisticated, digital logic will come to control, monitor, and interface with our bodies with increasing frequency. Exhibit 75.1 has a few examples of these devices and possible adverse events that may be caused by their failure or compromise. For example, at least one model of implantable cardioverter-defibrillators may be remotely reprogrammed or activated (!) with fairly inexpensive equipment.[5]

As computer systems that process data (often communicating this data wirelessly), these devices require traditional countermeasures to threats. The implications of compromise are much more severe than in other information systems, and may include the death of the patient. Consequently, controls that are strictly detective or respond after a compromise are less appropriate here.

Some interesting legal questions arise. Consider for a moment if an employee's implant required a reliable Wi-Fi connection to the Internet and a future interpretation of the United States' Americans With Disabilities Act required employers to provide such access as a reasonable accommodation. That could certainly force the affected organization toward support of a more broad BYOD policy. Consider now a different employee's *elective* implant, and whether that implant could be provided access, too.

The security lessons stretch beyond implantable medical devices (IMDs). Once surgically implanted, it may be infeasible, dangerous, or impossible to replace or remove an IMD. Consider the security posture required to deploy some nonimplanted IT system that might not be replaced or upgraded for five years (or ever[6]). I suspect that we can build effective systems with these constraints given that IT professionals have become used to supporting legacy systems for almost the entirety of the industry. When we assume that systems are legacy by design, can security, fault tolerance, and usability be maintained at a high level of service and performance? A redundant belt-and-suspenders approach may be more expensive initially, but the costs amortized over a very long life span can make sense.

75.4.3 Legal & Jurisdictional Issues. The law is not perfect, and neither are prescribed compliance requirements. While I do not intend to repeat material from

EXHIBIT 75.1 Potential Adverse Events in Various Implantable Medical Devices[a]

Device	Adverse Events
Pacemaker, implanted cardioverter defibrillator,[b] ventricular assist device[c]	Heart failure, tachycardia, bradycardia, arrhythmia[b and d]
Cochlear implant	Deafness, phantom sounds, distraction/confusion[e]
Prosthetic limb control system[f]	Injury, damage to prosthetic limb, inadvertent movement
Spinal cord stimulator[g]	Loss of pain relief, inappropriate stimulation
Sacral anterior root stimulator[h]	Infection from inability to void, inappropriate stimulation
Retinal prosthesis,[i] implanted contact lens, intraocular lens	Blindness, phantom images, distraction/confusion
Implanted infusion pump	Inappropriate dosage/timing
Brain-machine interface,[i] other neuroprosthesis[k]	Loss of consciousness, neural effects[l]
Responsive neurostimulator, other deep brain stimulator[m]	Inappropriate stimulation, failure to stimulate
Implanted monitor or sensor	Incorrect readings
Implanted radio frequency identification (RFID) tag[n]	Loss of privacy, data leakage
Implanted dynamic Light-Emitting Diode (LED) tattoo	Inappropriate display

[a] Jeremy A. Hansen and Nicole M. Hansen, "A Taxonomy of Vulnerabilities in Implantable Medical Devices," in *Proceedings of the Second Annual Workshop on Security and Privacy in Medical and Home-Care Systems* (SPIMACS '10) (New York: ACM, 2010): 13-20. DOI=10.1145/1866914.1866917

[b] M. Mirowski, M. M. Mower, W. S. Staewen, B. Tabatznik, and A. I. Mendeloff, "Standby Automatic Defibrillator: An Approach to Prevention of Sudden Coronary Death." *Arch Intern Med* 126, No. 1 (1970), 158–161.

[c] B. Glenville and D. Ross, "Ventricular Assist Devices," *British Medical Journal* 292 (February 1986): 361–362.

[d] S. Lee, K. Fu, T. Kohno, B. Ransford, and W. H. Maisel, Clinically Significant Magnetic Interference of Implanted Cardiac Devices by Portable Headphones," *Heart Rhythm* (October 2009).

[e] J. Robert, E. Morley, and E. J. Richter. Final Report on Measurements of Static Electricity Generated from Plastic Playground Slides, March 2006, www.access-board.gov/research/play-slides/report.htm (URL inactive).

[f] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz, "Cortical Control of a Prosthetic Arm for Self-Feeding," *Nature* 453, No. 7198 (May 2008): 1098–1101.

[g] R. B. North, D. H. Kidd, M. Zahurak, C. S. Jones, and D. M. Long, Spinal Cord Stimulation for Chronic, Intractable Pain: Experience over Two Decades," *Neurosurgery* 32, No. 3 (1993):384–395.

[h] G. S. Brindley, C. E. Polkey, and D. Rushton. "Sacral Anterior Root Stimulators for Bladder Control in Paraplegia," *Paraplegia* 6, No. 20 (1982):365–381.

[i] A. Y. Chow, V. Y. Chow, K. H. Packo, J. S. Pollack, G. A. Peyman, and R. Schuchard, "The Artificial Silicon Retina Microchip for the Treatment of Vision Loss from Retinitis Pigmentosa," *Archives of Ophthalmology* 122, No. 4 (2004):460–469.

[i] G. Santhanam, S. I. Ryu, B. M. Yu, A. Afshar, and K. V. Shenoy, "A High-Performance Brain-Computer Interface," *Nature* 442, No. 7099 (2006):195–198.

[k] Dong Song, Rosa H. M. Chan, Vasilis Z. Marmarelis, Robert E. Hampson, Sam A. Deadwyler, Theodore W. Berger, "2009 Special Issue: Nonlinear Modeling of Neural Population Dynamics for Hippocampal Prostheses," *Neural Networks* 22, No. 9 (November 2009), 1340–1351.

[l] T. Denning, Y. Matsuoka, and T. Kohno. "Neurosecurity: Security and Privacy for Neural Devices," *Neurosurgical Focus* 27, No. 1 (July 2009).

[m] F. T. Sun, M. J. Morrell, and R. E. Wharen, "Responsive Cortical Stimulation for the Treatment of Epilepsy," *Neurotherapeutics 5* (January 2008):68–74.

[n] J. Halamka, A. Juels, A. Stubblefield, and J. Westhues, The Security Implications of VeriChip Cloning," *Journal of the American Medical Informatics Association* 13, No. 6 (2006):601–607.

Chapter 64, "US Legal and Regulatory Security Issues," I do want to look toward some unresolved issues that will become more important. Even when laws are effective in achieving their stated goals, laws, like ethical guidelines for use, can rarely keep up with the breakneck pace of changes in technology. To choose a single example, emails older than 180 days on a service provider's server are considered abandoned by the Electronic Communications Privacy Act of 1986, which, despite multiple updates in the meantime, stands as current U.S. law at the time of this writing. Abandoned messages of this sort may be requested by any law enforcement agency without a warrant. This may have made sense in 1986, but with a significant portion of U.S. Internet users relying on Web- or cloud-based email, it no longer does.

Aside from the law's inability to keep up with technology, a few major issues come to mind:

- The increasing disconnect between policy makers and technology—policy makers are rarely knowledgeable enough to make decisions related to science and technology.
- Partly due to the previous point, and partly due to the increasing complexity of laws, rules, policies, and procedures, the increasing burden of regulatory compliance.
- Patchworks of changing state, national, and international laws can make compliance quite challenging, especially with the difficulty of pinning down data to a single location. Identical data on two different servers may fall into two jurisdictions, with different or conflicting regulations.
- The United States government's willingness to collect data on traffic (inadvertently or not) passing through its borders may violate privacy expectations of non-U.S. endpoints.

I admit that I don't have any suggestions here: tread carefully in this minefield, or get a job as a lawyer to get paid to hold the mine-detector for someone else!

75.4.4 Ubiquitous Technologies. With computers everywhere, the attack surface changes to everywhere. Where risks and the effects of exploitation were previously localized, these effects become equally pervasive and ubiquitous. Not only does each new device require security, but so do their future communications protocols. The devices' increasingly complex interactions in the Internet of Things make securing such interactions very difficult.[7] Not only does the number of targets increase, but so do the number of possible attack origins. Imagine a piece of malware that infects all the networked lightbulbs on the Internet and launches a distributed denial of service attack. Even if each bulb only produced a small bit of traffic, the number of bulbs might be staggering. The silver lining (if you can call it that) is that the number of easily compromised targets will also increase, making well-secured hosts less attractive.

While our information systems are becoming more numerous and complex, they are also becoming more autonomous. When data is processed faster than a human operator can react, computers often become decision makers with real-world impacts and consequences. Probably the best-known negative result of computerized decision making is the 2010 Flash Crash.[8] For whatever reason, algorithms monitoring and trading in the stock market decided that something (heuristically) bad was happening, and a feedback loop of sales and declining prices resulted in a drop of several percent in exchange-wide stock prices over about five minutes. Human administrators briefly

halted trading to stop the freefall, and prices recovered about fifteen minutes later. In 2013, rules to prevent such crashes went into effect,[9] but it is not clear that they will be effective at preventing future events, especially considering *mini flash crashes* seem to happen with great frequency.[10]

Among the lessons the Flash Crash offers, I would like to highlight three:

1. Increasingly fast and complex automated systems will require increasingly fast and complex automated security.
2. Automated systems must fail safely. Clearly defining the parameters of safety in large systems is as necessary as defining reasonable inputs to a function in a programming project.
3. We should not expect automated systems to work perfectly all the time. Human-operated fail-safes with useful and responsive interfaces are necessary for oversight.

To illustrate these lessons, consider the lessons in the context of a self-driving car:

1. Built into the car's driving algorithms are a multitude of collision avoidance, location, and speed sensors to keep passengers safe. Security systems that act as a supervisor must ensure that these systems are not disabled and should have the authority to revert those systems to a basic feature set, provide appropriate alerts, or initiate a response to failure.
2. A human passenger or remote human monitor must have the ability to seize control of a misbehaving vehicle.
3. Every failure condition must be handled and result in the safe stopping of the vehicle or other safe, correct behavior.

75.5 NEW DIRECTIONS. I explored a handful of technologies that may prove to be disruptive to our current information assurance paradigm, and before that, I suggested a skeleton for that future paradigm. Rather than wrapping up the chapter with restatement of those discussions or offering some witty sendoff, I will conclude with a not-fully-baked pitch for decentralized security that I hope will encourage you to continue to consider alternative and new directions in information assurance.

Empowering individual users of information systems is one way to deal with a fast-moving environment and ubiquitous attack surfaces. Despite each user's capacity to be a malicious actor, their supervisors should also trust them enough to be part of the solution. People can make mistakes, but people can also be terrific first responders to an incident.

Distributing responsibility to the users resembles peer-to-peer or crowdsourced applications, so each user can be responsible for his or her own job tasks and security tasks, the results of which may then be shared with other users, possibly without a centralized controller. Using these combined findings, we end up with ensemble forecasting[11] for security.

As the user interacts with the information systems, feedback about their interactions both automatically generated and user-triggered can shape the reputation of individual systems and users both in terms of performance and how well they support or provide the services they advertise. I like to think of this as an informal way to provide feedback with a *thumbs down* or *one out of five stars* in lieu of opening a formal trouble ticket. In this way, the network behaves much more like a recommender system with real-time

feedback and decision making. Since each system or user is dynamically assigned a reputation on the network, fuzzy authorization based on previous interactions with no a priori configuration is possible. When each device is a possibly malicious party, past behavior can be used as a predictor.

A related approach is to take advantage of existing social network relationships, such that individual users can be vetted by evaluating the links they share with previously trusted parties.[12] This mimics the trust and community interactions of the real world which, while imperfect, has evolved to be resilient to malicious and irresponsible parties, self-corrects, and requires no centralized authority. If we can strive for an imperfect by compromise-tolerant, disruption-tolerant, and flexible security posture, we will have a chance to ensure that technologies that move faster than we can keep up with will remain usable and effective at making our lives better. (That is what they're supposed to do, right?)

75.6 NOTES

1. Remember that according to the breach notification law SB1386: In California, everyone can hear you scream!
2. http://michaelnielsen.org/blog/quantum-computing-for-the-determined/
3. Aaronson, S. "The Limits of Quantum Computers." *Scientific American* 298(3):62-69, March 2008, http://www.scientificamerican.com/article.cfm?id=the-limits-of-quantum-computers
4. The first PQCrypto conference was held in 2006. See http://pqcrypto.org/
5. Halperin, Daniel, et al. "Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses." *Security and Privacy*, 2008. SP 2008. IEEE Symposium on. IEEE, 2008.
6. http://marsrovers.nasa.gov/overview/
7. http://www.wired.com/opinion/2013/01/securing-the-Internet-of-things/
8. http://www.bloomberg.com/news/2012-05-07/flash-crash-story-looks-more-like-a-fairy-tale.html
9. http://www.marketwatch.com/story/jitters-as-flash-crash-response-rules-take-effect-2013-04-08
10. http://money.cnn.com/2013/03/20/investing/mini-flash-crash/index.html
11. Michael Jahrer, Andreas Töscher, and Robert Legenstein. 2010. Combining predictions for accurate recommender systems. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining* (KDD '10). ACM, New York, NY, USA, 693-702. DOI= 10.1145/1835804.1835893 http://doi.acm.org/10.1145/1835804.1835893
12. Beach, Aaron, Mike Gartrell, and Richard Han. "Solutions to security and privacy issues in mobile social networking." *Computational Science and Engineering, 2009. CSE'09. International Conference on*. Vol. 4. IEEE, 2009.

INDEX

A posteriori testing, 10·8–10·9
A priori testing, 10·8
Abagnale, Frank, 2·5, 19·3
Abstract Syntax Notation 1 (ASN.1), 37·5, 37·18
Access control:
  access control entries (ACEs), 24·18–24·19
  access control list (ACL), 7·34, 16·14, 24·7–24·8, 26·5, 26·7, 32·6, 36·5
  access matrix, 24·7–24·8
  alarms. *See* Alarms
  audit trails. *See* Audit trail
  and authentication, 23·12, 23·14–15. *See also* Authentication
  breaching, 15·15–15·20
  bypass mechanism, 23·16
  card entry systems, 23·13
  data-oriented, 24·7–24·9, 66·9
  diagnostic utilities, 53·9
  discretionary access control list (DACL), 24·18–24·19
  distributed access control, 6·4
  e-commerce security services, 30·6
  file sharing access rights, 24·11–24·12
  file system, 36·5
  former employees, 13·2–13·3
  insurance for damage from, 60·14
  integrated systems, 23·15–23·16
  Internet-accessible systems, 30·38–30·39
  locks and door hardware. *See* Locks and door hardware
  log files, 53·16
  and malicious code, 16·3
  matrix model, 9·3–9·5
  methods of, 47·6
  monitoring, 53·20
  operations center, 47·5
  overview, 23·12
  and penetration techniques, 15·7
  physical, 16·3
  and piracy policy, 42·4
  portal machines, 23·16
  privileges, 23·12
  production programs and control data, 47·14–47·16
  proximity and touch cards, 23·13–23·14
  read-only access, 36·9
  radio-frequency identification (RFID), 23·12, 23·13
  surveillance systems. *See* Surveillance systems
  System Access Control List (SACL), 24·17–24·18
  user access to files and databases, 47·15–47·16
  user-oriented, 24·6–24·7
  visitor badges and log ins, 47·6
  Web servers, 30·31–30·32
  Websites, 15·27
  Windows 2000 security example, 24·14–24·20
  wireless LANs, 25·7–25·9, 33·7–33·8
  World Wide Web, 17·14
Access mask, 24·18–24·19
Access matrix, 24·7–24·8
Access token, 24·16–24·17
Accessibility, 4·20–4·21
Accidental threats, 13·5
Account permissions, 8·10–8·11
Acknowledgment numbers, 5·21
Active code, 26·16
Active Directory, 34·8
Active Server Pages (ASPs), 15·30, 18·13
Active taps, 6·5
ActiveX:
  buffer overflows, 39·13
  controls, 17·7–17·8, 17·12–17·14
  and e-commerce security, 21·8, 21·22
  malicious code, 16·10
  and mobile code, 17·3, 17·7–17·13, 21·8, 30·35
Ad-Aware, 21·9, 48·43
Addiction, 12·12, 48·15, 48·28, 48·38
Administrators:
  database administrator (DBA), 21·20
  information security administrators, 47·4–47·5
  passwords, access to by system administrators, 28·5–28·6
  professional accreditation and education, 63·27–63·28
  security administrators, 47·4–47·5, 63·25–63·28

Administrators (*Continued*)
system administrators, defenses against denial-of-service attacks, 18·25–18·26
system administrators, password access, 28·5–28·6
Adobe:
Acrobat, 7·27, 44·14
antipiracy programs, 42·5
Flash, 16·11, 41·10
Advance-fee fraud, 2·21, 16·13
Advanced Encryption Standard (AES), 7·39, 7·43–7·44, 34·9, 37·2
Advanced Television Committee Standard (ASTC), 42·9
Adware, 48·14–48·15, 48·42–48·43
Agentless scanning, 40·8
Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS), 11·35–11·39
Airborne particles, 22·21
AirMagnet, 33·45
Aircrack, 33·27, 33·45
Airopeek NX, 33·45
Airsnort, 33·45
AJAX, 17·3, 26·13, 26·16
Alarms:
duress alarms, 23·18
environmental, 23·18
fire, 23·27
intrusion alarms, 22·17
open-door alarms, 22·23, 23·10
overt and covert, 23·10
perimeter alarms, 23·17
premises-wide alerts, 23·27–23·28
silent, 23·27
Albert the saboteur, 2·4–2·5
Alerts, 53·22–53·23
Algorithms:
Blowfish, 7·21, 7·27
Data Encryption Algorithm (DEA), 7·20, 7·38
defined, 7·2
Diffie-Hellman, 7·23–7·24, 7·36–7·37, 37·16
encryption, 7·2–7·3, 7·8–7·9, 7·42–7·44, 37·16. *See also* Encryption
Grover's algorithm, 7·41–7·42
public key cryptosystems, 37·4–37·6, 37·16, 37·20, 37·23
risk assessment, 62·17–62·18
RSA algorithm, 7·25–7·27, 7·36–7·38, 37·16, 37·22–37·23
TKIP, 33·20
Secure Hash Algorithms (SHA), 34·9
Shor's algorithm, 7·41–7·42
All-hazard mitigation plan, 23·43–23·44
Allen, Paul, 1·9
Allowed path vulnerabilities, 21·7
Amazon.com, 48·25
America Invents Act of 2011 (AIA), 11·39
American Bar Association, *Digital Signature Guidelines,* 37·8–37·9
American Institute of Certified Public Accountants (AICPA), Security Audit Standards:
ISAE 3402, 54·19
SAS 70, 54·14
SSAE 16, 54·14–54·19
American National Standards Institute (ANSI), 6·23
ANSI/ISO/IEC 17024:2012, 74·8–74·9
ANSI X3.106-1983, 7·20
American Society for Industrial Security (ASIS), 74·11
Americans with Disabilities Act, 29·18
Anchoring, 19·13
Anomaly detection, 27·9–27·10
Anonymity:
benefits of, 70·17–70·18
and content filtering, 31·12
corporations, 70·26
cyberspace versus real world, 70·7–70·9
disadvantages of, 70·21–70·22
ethical principles, 70·24–70·26
families, 70·23–70·24
harassment and cyberbullying, 70·2–70·5
identity theft, 70·2
individuals, 70·23–70·24
Internet service providers, 70·26
model for identity in cyberspace, 70·27–70·28
overview, 70·1–70·2, 70·29–70·30
politics, 70·2
and privacy rights, 69·24–69·25, 70·18–70·20
pseudonymity, 70·17–70·18, 70·21–70·22
role of governments, 70·28
schools, 70·23–70·24
social psychology issues, 70·9–70·17. *See also* Social psychology
spam, 70·5–70·6
systems analysis, 70·22–70·23
terminology, 70·7–70·9
theory of nymity, 70·15
and Web monitoring, 31·12
Anonymizing remailers, 48·4
Anonymous, 2·28, 70·6
ANSI. *See* American National Standards Institute (ANSI)
Anti-Malware Testing Standards Organization (AMTSO), 41·5
Antimalware, 17·3–17·4, 26·16, 48·15, 66·10
Antisniffer tools, 25·4
Antispoofing, 26·10
Antispyware programs, 5·3, 48·15
Antivirus programs:
antivirus deployment 41·12–41·13
antivirus engines and databases, 41·6–41·7
automatic updates, 41·3
content filtering, 41·10–41·12
e-commerce security services, 30·7

and extent of malware, 41·1–41·2
issues, 41·3
for MacOS, 25·15
malicious code prevention, 16·12–16·14
personal computers, 5·3
policies and strategies, 41·13
scanners, 41·4–41·6
scanning methodologies, 41·7–41·10
terminology, 41·2–41·3
updating, importance of, 41·3
and VoIP, 34·7
and vulnerability assessments, 46·4
WildList, 41·5
Apple Computer, 1·8, 1·10, 2·15
Applets, 16·11, 17·3, 17·11–17·12, 21·8, 48·35
AppleTalk, 6·27–6·28, 25·15
Appliances, 26·7–26·8
Application isolation, 25·12
Application layer gateways (ALGs), 26·6, 34·7
Application servers, 21·8
Application service providers (ASPs), 30·44–30·45
Application tunneling, 31·11
Applications. *See also* Software
application-based monitoring for intrusion detection, 27·7
application layer gateways (ALGs), 26·6
application whitelisting technologies, 40·6–40·7
batch files 52·13
core layer, 5·10
data corruption 52·6–52·7
data integrity and validation, 52·13–52·16
database management subsystem, 52·7–52·8
databases, 52·5–52·6
design, 30·28
distributed databases, 52·8–52·12
input, 52·2–52·4
joint application development (JAD), 39·7
overview, 5·9, 52·1–52·2, 52·16
peer-to-peer (P2P), 5·26
rapid application development (RAD), 39·7
standards, 5·28–5·29
systems architecture and design, 52·2
user entitlements, 52·4–52·5
ARPANET (Advanced Research Projects Agency Network), 1·8
ASCII, 4·3
Association versus causality, 10·7–10·8
Association of Certified Fraud Examiners (ACFE), 74·11
Asymmetric attacks, 21·5
Asynchronous communications, 15·9
Asynchronous time, 4·11
Attacks:
asymmetric, 21·5
brute force attacks. *See* Brute force attacks
cafe latte attacks, 33·28–33·29
electromagnetic pulse attack (EMP), 22·20
email bombs, 18·6–18·8
endpoint attacks, 33·24
hotspot attacks, 33·27–33·28
insider, 13·2–13·8
Land attack, 18·10
layer 1 attacks, 33·23–33·24
overflow attacks, 15·29
security incident taxonomy, 8·12–8·16
Teardrop attacks, 18·10, 25·15, 26·5
WEP attacks, 33·29–33·32
Attitudes, 50·13–50·16
Audit trail, 23·12, 30·6, 53·11
Auditability, 26·29, 27·3–27·4
Auditor, 33·45
Audits:
audit controls, 19·18
audit file, 24·13
audit trail. *See* Audit trail
goals and responsibilities, 54·8
overview, 54·1–54·3
regulations for information security, 54·3–54·10. *See also* Gramm-Leach-Bliley Act (GLBA); International Organization for Standardization (ISO); Standards
standards, 54·10–54·19
technical frameworks, 54·26–54·28
Aureate/Radiate, 48·15
Authentication:
access control, 23·12, 23·14–23·15. *See also* Access control
biometric. *See* Biometric authentication
Captive Portal, 33·24–33·25
and Common Internet File System exploits, 36·10
costs of technologies, 28·18
cross-domain, 28·18
defined, 28·2
device authentication, VoIP, 34·8
e-commerce security services, 30·6
email, 20·25–20·26
and encryption, 7·5
extranets, 32·19
host site, 21·10–21·11
hotspot-based, 33·24–33·25, 33·27–3·28
IEEE standards. *See* IEEE 802 standards
issues, 28·17
MAC address, 33·24
and management frame spoofing, 33·26
and mobile code, 17·5
and operations security, 47·6
overview, 28·2, 28·18–28·20
person-to-computer versus computer-to-person, 28·2
personal identification number (PIN), 17·8, 23·14, 28·3, 28·9, 28·15
preliminary evaluation phase, security policy development, 66·5
principles of, 28·2–28·5
risk-based methods, 19·19

Authentication (*Continued*)
and security incident common language, 8·5–8·10
shared key (SKA), 33·8
token-based, 28·13–28·18. *See also* Tokens
two-factor, 25·6, 28·3, 28·9
user authentication, VoIP, 34·8
using mobile devices, 28·17–28·18
virtual private networks, 32·6
visual authentication, 23·14
and Wired Equivalent Privacy (WEP), 25·7–25·8
Authenticity:
e-commerce security services, 30·6
as source of loss, 3·2, 3·5–3·6, 3·8–3·12
and trust, 7·26–7·27
Authenticode, 17·6
Authority Revocation List (ARL), 37·19–37·20
Authorization:
defined, 28·2
e-commerce security services, 30·6
and sandboxes, 17·5
and security incident common language, 8·6, 8·13–8·15
Automatic updates, 16·13, 17·14
Automation:
computer security incident information, 56·16
AV-Test Institute, 51·22
Availability:
data, 24·3–24·4
extranet systems, 32·19
firewalls and gateway security devices, 26·20
hardware, 24·3
high-availability (HA) configuration, 26·24
software, 24·3
virtual private network (VPN), 32·13
Avalanche photodiode (APD), 6·10
Avatars, 48·30
Awareness programs:
accessibility, 49·26–49·27
analogies and examples, 49·22
audiences and messages, 49·14–49·15
awards and recognition, 49·36
behavior, security-related, 49·6–49·8
classroom training and clickers, 49·32–49·33
compliance statements, 49·30
contests, 49·34–49·36
conversational style, 49·22
credibility, 49·24–49·25
culture, 49·26–49·27
currency, 49·23–49·23
diversity, 49·26–49·27
evaluation, 49·38–49·41
goals, 49·13–49·14
Human Libraries, 49·37–49·38
images, 49·19
incentive prizes, 49·34–49·36
information security policy, 49·9–49·10
inspections and audits, 49·38
intranet Website, 49·29
key concepts, 49·3–49·5
management commitment, 49·10–49·11
messages and methods, pretesting and refining, 49·27–49·29
methods and frequency, 49·15–49·16
metrics, 49·8–49·9, 49·38–49·41
networked screen savers, 49·31
novelty and change, 49·21–49·22
overview, 49·2–49·3, 49·41
people as security problem, 49·4–49·5
penetration tests, 49·33–49·34
podcasts, 49·30
posters and digital signage, 49·31
publications, 49·31
resources, 49·11–49·12
sign-on messages, 49·31
social media and crowd sourcing, 49·30
social proof, 49·25–49·26
spaced repetition, 49·27
spear phishing exercises, 49·33–49·34
stories and empathy, 49·23
success factors, 49·8–49·16
surprise, 49·21
techniques, 49·18–49·28
terminology, 49·42–49·43
tools, 49·29–49·38
topics, 49·17–49·18
videos, 49·19–49·21, 49·30
visibility, audience appeal, and participation, 49·12–49·13
volunteer activities, 49·38
Web-based courses, 49·31–49·32

BabelFish, 31·12–31·13
Back door utilities, 15·26–15·27
Back Orifice, 15·22, 15·37, 21·4, 21·10
Back pocket file, 24·13
Backdoor.IRC.Snyd.A, 17·5
Background checks, 16·12, 45·2–45·3
Backoff schemes, 6·13–6·14
Backout and recovery, 47·7– 47·9
BackTrack, 33·45
Backup generators, 23·30–23·36
Backups:
data. *See* Data backups
and data losses, 3·4, 3·7, 3·15
data storage, 36·5–36·6
encryption, 36·6
file creation, 52·13
hardware, 4·18–4·19
and new versions of software, 47·7–47·9
personnel, 4·18
plans for, 4·17–4·18
power, 4·19–4·20
security policies for, 66·10–66·11
testing, 4·20
Backward learning, 6·25
Bandwidth consumption, 18·8–18·10

Banner ads, 21·8
Baseband bus, 6·6
Bastille Linux, 32·17
Batch files 52·13
BD-R discs, 4·9
BD-RE discs, 4·9
Beacon frame spoofing, 33·26–33·27
Behavior. *See also* Psychology; Social psychology
  explanations of, 50·7
  group behavior, 50·20–50·21
  management, 56·31
  professionalism, 56·26–56·27
  prosocial, 50·16–50·17
  social pressure and behavior change, 50·18
Beliefs and attitudes, 50·13–50·16
Bell-LaPadula model, 9·2, 9·9–9·12, 9·18–9·19
Berkeley Software Distribution (BSD) License, 11·34
Berne Convention, 11·35, 11·37
Biba's strict integrity policy model, 9·2, 9·9, 9·12–9·14, 9·18–9·19
Binary code analysis, 38·16–38·17
Binary design, 4·2–4·4
Bind, 21·12
Biological and chemical warfare, 23·42
Biometric authentication:
  and authentication principles, 28·2, 28·4
  biometrics, 29·4–29·5
  costs, 29·19
  crossover error rate (equal error rate), 29·16
  disadvantages and concerns, 29·17–29·22
  and double enrollment prevention, 29·7
  dynamic biometrics, 28·5
  encryption, 29·19, 29·21
  enrollment, 29·7–29·8, 29·17
  facial scans, 23·15, 29·11–29·12, 29·16, 29·19
  facilities access (physical access), 29·7–29·8
  failure to enroll, 29·17
  false accepts, 29·16
  false rejects, 29·16
  finger scans, 23·14–23·15, 29·9–29·11, 29·16, 29·18, 29·20
  fraud, 29·19
  government use of, 29·22
  hand geometry scan, 29·12–29·13, 29·18
  history of, 29·2–29·4
  iris scan, 23·15, 29·13–29·14, 29·18
  keystroke scans, 23·15, 29·15
  overview, 28·20, 29·2–29·4, 29·24, 29·26
  privacy issues, 29·17, 29·19, 29·20–29·22
  public identification systems, 29·8
  retinal scanning, 29·15
  security (logical access), 29·7
  signature scans, 29·15
  static biometrics, 28·4
  technologies, comparison chart, 29·25
  templates for data acquisition, 29·8–29·9
  time and attendance systems, 29·24
  trends, 29·22–29·24
  verification of identity, 29·5–29·7
  voice recognition, 23·15, 29·14–29·15, 29·19
Biometric scanners, 23·14, 23·15–23·16
Bitlocker Drive Encryption, 25·12
BitTorrent, 35·8
Black box testing, 38·17–38·18
Black boxes, 69·14
Block lists, 20·20–20·21, 31·9
Blogs, 48·5, 48·31
Blu-ray, 57·8
Bogus network addresses (BOGONs), 16·13, 26·10
Boot sector viruses, 16·4. *See also* Viruses
Border Gateway Protocol (BGP), 5·23, 5·27
Bot herders, 16·11
Bot Roast II, 17·13
Botnets, 2·21, 15·31–15·32, 55·13
Bots, 11·28, 16·8, 17·13, 18·7
Boundary condition violations, 38·11, 39·9
Brainbench basic security certifications, 74·11–74·12
Brick walls, 17·5
Bridges, 6·25–6·26, 37·12
Broadband access lines, 5·4
Broadband bus, 6·6
Broadband Technology Advisory Group (BBTAG), 6·18
Broadband Wireless Access (BBWA) Working Group, 6·18
Brownouts. *See* Power failures and disturbances
Brute force attacks:
  cryptanalysis, 7·9–7·11, 7·17, 21·12, 37·22
  penetration techniques, 15·15
  and VoIP theft of service, 34·6
Buffer overflow:
  application server program, 21·20
  boundary violations, 38·10
  denial-of-service (DoS) attacks, 18·8
  and e-commerce, 21·12
  and IM applications, 35·11
  and network file systems, 36·9
  and system penetration, 15·24–15·25, 15·28
  and testing software, 39·13
  UNIX, 25·14
  Web and mail servers, 21·4
Bugs:
  commercial off-the-shelf software, 21·13
  debugging, 39·5, 39·19–39·20
  seeding, 39·15
  software, 5·9
  tracking and removal, 39·17
  and wiretapping, 23·41–23·42. *See also* Wiretapping
Build Security In (BSI), 38·15
Building Security in Maturity Model (BSIMM), 51·8, 51·14
Burma (Myanmar), censorship in, 72·11
Bus topology, 6·6

Bus-wired ring, 6·8
Business continuity planning (BCP):
  business impact analysis, 58·14–58·29. *See also* Business impact analysis (BIA)
  and computer security incident response team, 56·7–56·8
  and corporate mission, 58·10–58·11
  cost justification, 58·29–58·34
  emergency issues, 58·13–58·14
  goals, 58·8–58·14
  generalized cost consequence (GCC) model, 58·31–58·34
  overview, 58·1–58·2
  presentation of plan, 58·34–58·36
  purpose of, 58·2
  quantitative risk model, 58·29–58·31
  recovery phases, 58·13
  recovery scenarios, 58·6–58·7
  and risk, 58·2–58·4
  scope of plan, 58·9–58·10
  types of disasters, 58·4–58·6
  validating goals, 58·11–58·12
Business impact analysis (BIA):
  criticality of functions, 58·20–58·22
  cross-department functions, 58·26–58·27
  departments and functions, 58·18–58·19, 59·19–59·20
  function category, 58·23
  functions, 58·17–58·18
  interview process, 58·15–58·18
  key persons and key alternate, 58·19–58·20
  matrix analysis, 58·24–58·29
  operational impact, 58·22–58·23
  overview, 58·14
  ranking factor, 58·27–58·29
  scope of, 58·15–58·16
  survival time, 58·20
  system elements, 58·24
Business resumption planning (BRP), 66·11
Business-to-business (B2B) transactions, 21·6–21·8, 30·13–30·17
Business-to-customer (B2C) transactions, 30·3–30·4, 30·9–30·13
BUTTsniffer, 25·4
Bytes, 4·3–4·4, 5·2

C, 21·14, 21·20
C++, 21·14, 47·3
Cable carriers, 4·15–4·16
Cable Communications Policy Act of 1984, 69·21
Cable Modem Working Group, 6·18
Cables:
  coaxial cable (coax), 6·2, 6·8–6·9, 6·13
  and threats to information infrastructure, 22·17–22·19, 23·26
Cache files, 15·18
Caching, 26·16
Caching services, 31·13
Caesar cipher, 7·8, 7·7–7·11
Cafe latte attacks, 33·28–33·29
Calculation errors, 38·11, 39·9
Cameras. *See* Surveillance systems
CAN-SPAM Act of 2003, 16·11, 20·15, 20·26–20·27
Cantennas, 33·21
Canter, Laurence A., 2·20–2·21
Capability matrix, 45·10–45·11
Capability Maturity Model (CMM), 51·7, 51·11–51·12
Capability tickets, 24·8–24·9
Card entry systems, 23·13
Carrier sense multiple access with collision avoidance (CSMA/CA), 6·14
Carrier sense multiple access with collision detection (CSMA/CD), 6·13–6·14, 6·19–6·20, 6·25
Cascade virus, 2·11
Catastrophic events. *See* Physical threats
Cathode-ray terminals (CRTs), 5·12
Causality versus association, 10·7–10·8
CD-R, 4·9
CD-ROM, 4·9
CD-RW, 4·9
CDc, 15·39
Cellular phones, 15·12, 21·8, 35·16. *See also* Short message service (SMS)
Censorship. *See also* Content filtering
  First Amendment rights, 72·12–71·21
  government intervention, 72·2–72·3
  hate speech, 72·3–72·4
  objectional material, defined, 72·3
  overview, 72·22
  parental involvement/responsibility, 72·21–72·22
  politics and culture, 72·4–72·12
  sexually oriented speech, 72·3
Central processing unit (CPU) log record, 53·14
Certificate Policy (CP), 37·8–37·9, 37·17
Certificate Revocation List (CRL), 7·35, 37·6, 37·13, 37·16–37·23
Certification Authority (CA), 17·6, 21·10–21·11, 33·37, 37·5–37·22
Certification practice statement (CPS), 37·8–37·9
Certified Wireless Network Professional (CWNP), 74·16
Certified Wireless Security Professional (CWSP), 74·16
CGI. *See* Common Gateway Interface (CGI)
CGI/PHP (Hypertext Processing), 21·13–21·15
Chain letters, 48·10–48·11, 48·15
Challenge-handshake authentication protocol (CHAP), 25·9–25·10
Change:
  incremental change, 50·19–50·20
  management, 26·21–26·22, 39·16–39·18, 53·24–53·25
  operations staff responsibilities, 47·7–47·9
  social psychology. *See* Social psychology

Chaos Computer Club (CCC), 2·23, 17·8–17·9
Chargeback systems, 53·23
Check digits, 47·17
Check Mark program, 51·21
Checklist for hardware security, 4·25–4·27
Checksums, 7·4, 39·20, 41·9, 46·4, 47·9, 53·16
Chief information security officer (CISO):
  as change agent, 65·1–65·3
  alliance with counsel, 65·16
  culture of security, 65·15–65·16
  executive visibility, 65·13
  governance and accountability, 65·9–65·11
  insecurity of systems, 65·5
  metrics, 65·13
  monitoring, 65·13
  organizational structure, 65·17–65·18
  overview, 65·13–65·14, 65·18–65·19
  partnership with internal audit, 65·16–65·17
  recommendations for, 65·14–65·18
  relationship with IT, 65·17
  reliance on digital information, 65·4–65·5
  reporting, 65·13
  responsibilities, 63·2–63·4, 65·11–65·13, 65·18
  standard of care, 65·6–65·7, 65·9
  as strategist, 65·3–65·4, 65·6
  world trends, 65·5
Child Internet Protection Act (CIPA), 31·3
Child Online Privacy Protection Act, 48·15
Child pornography, 48·13, 61·5
Children:
  and cybersex sites, 48·28–48·29
  and Internet pornography, 48·36
  and plagiarism, 48·32
  protecting, recommendations for, 48·37
  and video games, 48·29–48·30, 48·40
Children's Online Privacy Protection Act of 1998 (COPPA), 11·29, 30·20, 69·18–69·19
China:
  censorship in, 72·5–72·7
  and piracy, 11·20
  reliability of Chinese-manufactured computer components, 2·14
  technological attacks on defense and R&D facilities, 16·3
Chinese Wall model (Brewer-Nash model), 9·2, 9·17–9·18
Chipping code, 25·7
Choice-supportive bias, 19·12
Christmas Tree worm, 2·16, 18·3, 18·7
Ciphers, 4·17, 7·2, 7·6–7·16
Ciphertext, 7·2–7·3, 37·2–37·3, 37·2–37·3
Citibank, 2·10
Cladding, 6·11
Clark-Wilson model, 9·2, 9·9, 9·14–9·16, 9·18–9·19
Clearing Houe Interbank Payment System (CHIPS), 14·14
Client puzzle protocol, 18·29
Clinical Information Systems Security model, 9·18
Closed circuit television (CCTV) surveillance cameras, 23·17
Cloud computing, 1·13, 47·18–47·20, 53·28–53·29
Cloud Data Management Interface (CDMI), 57·17
Cloud Security Alliance (CSA), 51·11, 51·17
Cloudfogger, 35·4
Coalition Against Unsolicited Commercial Email (CAUCE), 20·14, 20·18
Coaxial cable (coax), 6·2, 6·8–6·9, 6·13
COBOL, 47·3
Code Red Worm, 18·23–18·24
Codes and coding. *See also* Encryption assurance tools
  assurance tools, 38·15–38·20
  best practices and guidelines, 38·7–38·9, 38·20
  and binary design, 4·2–4·4
  buffer overflow vulnerabilities, 38·15. *See also* Buffer overflow
  decoding and debugging phase of software development, 39·5
  design phase, 38·6
  digital signatures, 38·8, 47·9
  due diligence, 38·4–38·5
  errors, 16·13, 21·7, 38·10–38·15
  languages, 38·9–38·10
  operating system considerations, 38·6–38·7
  policy and management issues, 38·1–38·5
  regulatory compliance issues, 38·5
  requirements analysis, 38·5–38·6
  secure code, overview, 38·1
  security integration, 38·4
  self-checking codes, 4·5
  self-replicating code, 16·2
  signed code, 17·6–17·11
  software errors, 38·10–38·15, 39·8–39·12
  software total quality management, 38·2–38·4
  source code, 47·3, 47·7, 47·9
  Trojan code, 20·1–20·3, 20·31–20·35
Codes of conduct, 69·29
Cognitive bias, 19·12–19·13
Collaboration tools, 35·3–35·4, 35·18–35·20
Coloured Petri Nets (CPNets), 55·24
Commercial and consumer Web transactions, 30·43–30·44
Commercial off-the-shelf (COTS) software, 14·3–14·4, 17·4, 21·13, 21·22, 47·9
Commercial Product Evaluation Program, 24·14
Committee on National Security Systems (CNSS), 74·9–74·10
Committee of Sponsoring Organizations (COSO) of the Treadway Commission, 54·18, 54·21–54·22
Common Criteria (CC), 51·3–51·4, 51·26–51·28, 71·47

Common Criteria Evaluation and Validation Scheme (CCEVS), 54·4
Common Gateway Interface (CGI), 15·27–15·28, 21·2–21·4, 21·13–21·15, 21·17
Common Internet File System (CIFS), 36·4, 36·9–36·10
Common language for security incidents, 8·1–8·20, 55·14
Common Object Model (COM), 21·13
Common Object Request Broker Architecture (CORBA), 21·13
Communications:
  data, 4·13–4·16, 4·23–4·24
  intercepting, 15·8–15·15
Communications Assistance for Law Enforcement Act (CALEA), 34·3
Communications Decency Act (CDA), 72·18
Compact disc read-only memory (CD-ROM), 4·9
Compact discs (CDs), 36·1
Compartmentalization, 30·37–30·38
Compatible Time Sharing System (CTSS), 1·7
Component-based software (CBS), 21·13–21·14
CompTIA, 74·14–74·15
Computer crime:
  data diddling, 2·9–2·10
  data theft, 4·22
  denial of service, 2·21–2·22
  detection, 10·2
  embezzlement, 45·6
  and employees, 45·1–45·2. *See also* Employees
  extortion, 2·12
  hacker underground, 2·22–2·23
  historical records, studying, 2·2
  identity theft, 2·7–2·8, 70·2
  impersonation, 2·5–2·8
  industrial espionage, 2·29–2·31
  law enforcement involvement, security incidents, 56·11
  logic bombs, 2·11, 16·4–16·5, 45·9
  and malicious code, 16·3, 16·14
  online gambling, 48·26–48·27, 48·41–48·42
  overview, 2·2–2·3
  pedophiles, 48·13
  phone phreaking, 2·8
  Ponzi schemes, 48·10–48·11
  reporting, 10·2
  research methodology, 10·3–10·11
  sabotage, 2·3–2·5
  salami fraud, 2·10–2·11
  social engineering. *See* Social engineering
  spam, 2·20–2·21
  statistics on, limitations, 10·3–10·4
  Trojan horses. *See* Trojan horses
  unauthorized security probes, 45·11–45·12
  workplace issues, 48·33
  worms and viruses, notorious, 2·15–2·20
Computer criminals. *See also* Computer crime; Cyberinvestigation
  aggressive behaviors, 12·4–12·9
  anonymity, 12·4–12·5, 12·7, 12·16–12·17
  antisocial personality disorder, 12·9–12·10, 12·18
  and Asperger syndrome, 12·11–12·12
  categories of, 8·16
  classifications of, 12·17–12·24
  computer addiction, 12·12–12·13
  cyberterrorism, 12·20
  cyberterrorists, 12·20
  dangerous/malicious insiders (DI/MI), 12·21–12·22
  deindividuation, 12·6–12·7
  ethical immaturity, 12·14–12·16
  five-factor model of personality, 12·10–12·11
  hackers, 12·2–12·4, 12·16–12·17
  motivation, 12·3–12·4, 12·20–12·21, 55·17–55·20
  narcissistic personality disorder, 12·9–12·10
  overview, 12·1–12·3
  personality disorders, 12·9–12·10
  profiling, 55·17–55·20
  recommendations, 12·24–12·25
  and social context, lack of, 12·5–12·6
  social identity model of deindividuation effects, 12·7–12·8
  social learning theory, 12·8–12·9
  social presence, 12·6
  social psychology. *See* Social psychology
  taxonomy, 12·19–12·20
  victim-blaming, 12·4–12·5, 12·22
  virus creators, 12·22–12·24
Computer Emergency Response Team Coordination Center (CERT/CC), 8·2–8·3, 15·14, 18·3, 18·15, 56·4–56·5, 74·14
  assistance during security incident, 56·11
  reporting incidents to, 56·35–56·36
  security improvement modules, 44·5–44·6
Computer Fraud and Abuse Act (CFAA), 11·27–11·33, 16·9, 16·11
Computer languages, 38·9–38·10, 47·3
Computer program, defined, 47·3
Computer Security Incident Handling Guide, 56·34
Computer security incident response team (CSIRT):
  baselines, 56·12, 56·13–56·14
  burnout, 56·29–56·30
  common services, 56·5–56·6
  conferences, attending and speaking at, 56·36
  contacts, establishing list of, 56·10–56·12
  continuous process improvement, 56·7, 56·34–56·35
  description, 56·3
  documentation, 56·17–56·21, 56·24
  emergency response, 56·22–56·26
  establishing, 56·7–56·10

forensic awareness, 56·32
functions of, 56·5
helpdesk, role of, 56·21, 56·27
history and background, 56·4–56·6
in-house versus outsourcing, 56·6–56·7
law enforcement involvement, 56·11, 56·24–56·25
legal staff involvement, 56·13
managing, 56·26–56·32
mission and charter, 56·7–56·8
outside agencies, interaction with, 56·10–56·12
overview, 56·2–56·3, 56·36–56·37
policies and procedures, 56·7–56·10
postincident activities, 56·32–56·36
postincident analysis, 56·33–56·34
public affairs, role of, 56·32
purpose of, 56·3–56·4
reporting security incidents, 56·35–56·36
role of, 56·3–56·4
selection of team, 56·12–56·13
shiftwork, 56·31–56·32
and social engineering, 56·29
technical expertise, 56·12–56·13, 56·16
telephone hotline, 56·21–56·22
tracking incidents, 56·17–56·21
training, 56·16–56·17
triage, 56·13, 56·14–56·16, 56·27–56·29
types of teams, 56·6–56·7
Computer Security Institute (CSI), 56·36
Computers:
internal clocks and time issues, 4·10–4·11
LAN components, 6·2–6·3
laptops, 1·19, 36·12–36·13
large-scale, 1·4–1·5
medium-size, 1·5–1·6
personal computers (PCs), 1·8–1·10, 4·20–4·25, 5·3
small-scale, 1·6
system assets, 24·3–24·4
Computers and Risk, 1·13
Conduit, 23·26
Confidence limits, 10·6–10·7
Confidentiality. *See also* Privacy
and disclosure of information, 3·13, 3·16
employee nondisclosure agreements, 45·16
and encryption, 7·4. *See also* Encryption
as source of loss, 3·2, 3·6–3·14
and WLANs, 33·8
Configuration:
endpoint supplicant configurations, 33·36–33·37
firewalls and gateway security devices, 26·23–26·24
protection, 53·25
secure authentication server configuration, 33·37
Web servers, 21·17, 21·22
Confirmation bias, 19·12
Conformity, compliance, and obedience, 19·13–19·14
Consequentialism (teleology), 70·25
Construction concerns:
and clean electrical power, 23·28–23·30
confidential design details, 23·6–23·7
electrical power, 23·28–23·36
electrical standards, 23·8
emergency power, 23·30–23·36
equipment cabinets, 23·10
equipment rooms, 23·25–23·27
facilities design, 23·21–23·41
protective construction, 23·24–23·27
telecommunications standards, 23·8
violence prevention and mitigation, 23·7–23·8
Content filtering. *See also* Censorship
and antivirus technology, 41·10–41·12
pornography, 48·36
Contention network, 6·13–6·14
Contingency tables, 10·7
Continuous process improvement, 56·7, 56·34–56·35
Control (command) structure errors, 38·14, 39·11–39·12
Control groups, 10·8
Control mode, 24·9
Control Objectives for Information and Related Technology (COBIT), 44·5, 53·4, 54·22–54·24, 65·8
Controls:
ActiveX, 17·7–17·8, 17·12–17·14. *See also* ActiveX
validation controls, 52·13–52·16
Controlling the Assault of Non-Solicited Pornography and Marketing Act of 2003, 11·29
Cookie poisoning, 21·18
Cookies, 19·21, 21·8–21·9, 21·18–21·19, 48·25, 69·25
COPS vulnerability assessment system, 46·3
Copying, as means of information loss, 3·2, 3·16
Copyright law:
copyright ownership, 11·9
database protection, 11·21
derivative works, 11·12
Digital Millennium Copyright Act, 11·14–11·18, 11·22–11·23
fair use exception, 11·10, 11·16–11·17, 11·21–11·23
first sale doctrine, 11·9–11·10
formulas, 11·10
hyperlinks, 11·23
infringement, 1·20, 11·13–11·14
interfaces, 11·11
international law, 11·35–11·39
and licensing, 11·9–11·10. *See also* Licenses and licensing
"look and feel" of software, 11·10–11·11
overview, 11·8, 42·2

Copyright law (*Continued*)
registration of copyright, 11·9
remedies for infringement, 11·13–11·14
remedies for violation of DMCA, 11·18
and reverse engineering, 11·11, 11·16–11·17
Semiconductor Chip Protection Act, 11·13
transformative uses, 11·12, 11·21–11·24
and TRIPS, 11·37
watermarks, use of, 42·11–42·12
works for hire, 11·9
Core layers, 5·10
Corporate espionage, 14·15–14·16
Cost-benefit analysis, information infrastructure protection, 23·44–23·45
Costs:
authentication technologies, 28·18
biometric authentication, 29·19
data backups, 57·32–57·33
denial-of-service attacks, 18·4–18·5
physical threats, 22·5, 22·13–22·14
Public Key Infrastructure (PKI), 37·27
spam, 20·10–20·12
Council for Responsible Email (CRE), 20·16
Counter Mode with Cipher-Block Chaining Message Authentication Code Protocol (CCMP), 33·8, 33·10, 33·16–33·20
CoWPAtty, 33·45
C.P. Direct case, 20·8–20·9
Crack, 15·25, 25·9
Credit card fraud, 2·6–2·7
Credit card transactions, 21·8
Crime and criminals, reporting, 22·3
Criminal liability:
exposure of consumers' PII, 39·13
and litigation issues, 30·43
Cross-certification, 37·14
Cross-domain authentication, 28·18
Cross-site request forgeries (CSRF), 21·19
Cross-site scripting (XSS), 16·6, 21·19, 32·18
CryptoCard, 28·17
Cryptographic Algorithm Validation Program (CAVP), 51·23
Cryptographic Module Validation Program (CMVP), 51·23
Cryptographic viruses, 16·7
Cryptography, 4·16–4·17, 7·3–7·16, 14·23, 18·20, 37·2. *See also* Encryption
Cult of the Dead Cow (cDc), 2·23–2·24, 12·20
Customers:
business-to-customer security services, 30·3–30·4, 30·9–30·13
defined under Gramm-Leach-Bliley Act (GLBA), 64·8
loss of, 30·24
CyberCash, 21·10
Cyberinvestigation:
analysis of individual events, 55·10, 55·16
attack-tree analysis, 55·24
correlation, 55·11, 55·16–55·17
cyberforensics compared, 55·2
deconfliction of events, 55·11, 55·16
describing attacks, 55·14–55·15
end-to-end digital investigation (EEDI), 55·2, 55·9–55·17
end-to-end process, 55·9–55·12
evidence, 55·10, 55·12, 55·16–55·17
framework for, 55·2–55·9, 55·12–55·17
intrusion process, 55·13–55·14
investigative narrative, 55·13
means, 55·17–55·18, 55·20
modeling, 55·24–55·25
motive, 55·17–55·20
normalizing events, 55·11, 55·16
opportunity, 55·17–55·18, 55·20
overview, 55·1–55·9, 55·25
Rogers taxonomy, 55·2–55·3
strategic campaigns and tactical attacks, 55·15–55·16
threat agents, 55·17–55·20
timeline analysis, 55·12, 55·17
tools for, 55·21–55·25
CyberSecurity Institute, 74·13–74·14
Cybersecurity professionals, 74·1–74·8. *See also* Information Assurance (IA)
CyberSecurity Watch Survey, 13·3
Cybersex, 48·28–48·29
Cyberspace, defined, 70·7
Cyberterrorism, 12·20–12·21, 14·15–14·16. *See also* Information warfare (IW)
Cyberwar. *See* Information warfare (IW)
Cyclical redundancy check (CRC), 41·9, 53·16
Cylink, 7·27

Daemons, 15·27, 15·28, 18·17–18·18, 25·14, 26·18
DARPA. *See* Defense Advanced Research Projects Agency (DARPA)
Dashboards, 53·2–53·3, 53·24
Data:
access to, 47·14–47·16, 66·9. *See also* Access control
aggregation and reduction, 53·21–53·22
backups. *See* Data backups
classification, 30·7, 66·5, 67·1–67·16
collection, 66·3–66·4
communications, 4·13–4·16, 4·23–4·24
confidentiality, 33·8
corruption, 39·19–39·20, 52·6–52·7
data life cycle management (DLM), 57·24–57·27
destruction, 30·7, 66·10–66·11
diddling, 2·9–2·10, 3·15
files, defined, 47·4
integrity, 24·4, 47·17, 52·13–52·16
leakage, 15·7–15·8, 26·15
media longevity, 57·26
protection, 47·14–47·16
retention, 30·7, 57·25

rotation, 57·25–57·26
scavenging, 15·18
security, 24·3–24·4
storage. *See* Data storage
technology changes, 57·26–57·27
test data, 47·15
theft, 4·22
validation, 47·16–47·18, 52·13–52·16
Data backups:
archives, 57·2, 57·14–57·15
Blu-ray discs, 57·8
bootable versions, 57·15–57·16
cloud, 57·2, 57·16–57·18
and computer systems, 57·21–57·24
content-addressed storage (CAS), 57·14
costs of, 57·32–57·33
data storage capacities, 57·2
delta, 57·20
differential, 57·19–57·20
and disaster recovery, 59·13
disk mirroring, 57·3–57·4
disposal of backup media, 57·29–57·32
double, 57·2
DVD, 57·8
environmental protection, 57·27
external hard disk drives, 57·7–57·8
flash drives, 57·12
frequency of, 57·33–57·34
full, 57·19
hierarchical storage systems, 57·3
holographic disks, 57·8–57·9
incremental, 57·2, 57·20
indexing, 57·14–57·15
known-good boot media, 57·24
labeling, 57·13–57·14
laptops, 57·23
logging, 57·6–57·7
mainframes, 57·21–57·22
Millipede, 57·12
mobile devices, 57·23–57·24
need for, 57·1–57·3
network-attached storage (NAS), 57·4
offsite protection, 57·28–57·29
onsite protection, 57·27–57·28
optical storage, 57·8–57·9
policies, 57·35
recovery, 57·6–57·7
recovery point objectives (RPO), 57·3
redundant array of independent disks (RAID), 57·4–57·5
registry backups, 57·15
removable hard disk drives, 57·8
removable media, 57·7–57·8
scavenging, 57·30
servers, 57·22
software for, 57·6–57·7
storage area network (SAN), 57·4
strategies, 57·18–57·24
tape cartridge systems, 57·9–57·10
terminology, 57·2
testing, 57·24
Virtual Tape Library (VTL), 57·10–57·12
workstations, 57·4, 57·6, 57·22
Data classification schemas, examples of, 67·9
Data diddling, 2·9–2·10, 3·15
Data Encryption Algorithm (DEA), 7·20, 7·38
Data Encryption Standard (DES), 7·2, 7·16, 7·17–7·21, 25·11, 37·2
Data Execution Prevention, 25·12
Data life cycle management (DLM), 57·24–57·27
Data loss prevention (DLP), 13·8–13·10, 26·15–26·16
Data Protection Directive (EU), 11·29, 11·37
Data storage. *See also* Storage media
backup security, 36·5–36·6
best practices, 36·2–36·3
database encryption, 36·13–36·15
direct attached storage (DAS), 36·4
disposal of data, 36·15–36·16
encryption, 36·10–36·15
fibre channel threats, 36·7, 37·8
file system access controls, 36·5
in-band management, 36·4–36·5
leakage using cloud storage, 13·9
and management interfaces, 36·6–36·7
network attached storage (NAS), 36·4
network file system threats, 36·8–36·9
nonvolatile media, 36·1
out-of-band management, 36·4–36·5
overview, 36·1–36·2, 36·16
restore system controls, 36·5–36·6
sanitizing electronic media, guidelines for, 36·15–36·16
security basics, 36·2
smart phone encryption, 36·15
storage area network (SAN), 36·3–36·5
Database administrator (DBA), 21·20
Databases:
and application control, 52·5–52·6
attacks, 21·7–21·8
backout, 52·12
backup files and system logs, 52·11–52·12
controls to maintain integrity, 52·9
copyright protection, 11·21
database management subsystem, 52·7–52·8
deadlocks, 52·10
distributed databases, 52·8–52·12
encryption, 36·13–36·15
lock on update, 52·9–52·10
privacy issues, 21·11
recovery and restart, 52·12
roll-forward recovery 52·12
security, 21·19–21·20
two-phase commit, 52·10–52·11
unconditional vs. conditional locking, 52·10
Datacomp hardware Trojan, 2·13
Daubert v. Merrell Dow Pharmaceuticals, 73·2–73·4. *See also* Expert witnesses

de Guzman, Onel, 2·20
Decryption, 7·2–7·3, 7·9–7·10, 7·26, 7·37
Deducibility security, 9·19
Defacement of Web pages, 15·28, 15·30
DefCon, 15·39
Defective by Design, 42·16–42·17
Defense Advanced Research Projects Agency (DARPA), 1·13, 5·7, 7·42, 56·5
Defense in depth, 16·12, 16·14, 21·5, 36·6–36·7
Defense Information Systems Agency (DISA):
  computer security incident training, 56·8–56·9, 56·13, 56·15, 56·26, 56·28–56·30
  Security Technical Implementation Guidelines (DISA-STIG), 54·4
Delta CRL, 37·20–37·21
Demand Priority Working Group, 6·18
Demilitarized zone (DMZ), 30·29, 30·35
Demon (war) dialing, 15·16, 21·12
Denial-of-service (DoS) attacks. *See also* Distributed denial-of-service (DDoS) attacks
  accidental, 30·39
  Active Server Pages (ASP) attacks, 18·13
  arnudp, 18·13
  and autoforwarding e-mail, 48·22–48·23
  bandwidth consumption, 18·8–18·10
  billion laughs,
  black fax attack, 18·13
  boink, 18·13
  bonk, 18·13
  buffer overflow, 18·8
  cancelbot, 18·13
  Code Red Worm, 18·23–18·24
  costs of, 18·4–18·5
  decline of, 55·14
  defined, 8·15
  destructive devices, 18·6
  Domain Name System attacks, 5·25, 18·10–18·11
  email bombs, 18·6–18·8
  history of, 2·21–2·22, 18·2–18·4
  and information warfare, 14·21–14·22
  and instant messaging, 35·10–35·11
  and Java, 18·13
  and LANs, 25·5
  and loss of information, 3·15
  management issues, 18·30–18·31
  NIMDA, 18·24–18·25
  overview, 18·2
  PHP attacks, 18·13
  ping, 18·9
  Ping of Death, 18·8, 25·15
  preventing and responding to, 18·14
  resource starvation, 18·12–18·13
  router attacks, 18·13
  routing and Domain Name System attacks, 18·10–18·11
  Slowloris attack, 18·12
  SYN flooding, 18·11–18·12
  types of, 18·6
  using exploitable software, 18·23–18·25
  Website attacks, 15·27
  Windows NT servers, 18·4
  and Windows XP, 25·11
Dense wave division multiplexing (DWDM), 6·10
Department of Defense (DoD):
  Common Criteria, 71·47
  and Common Language Project, 8·3
  Computer Security Center, 24·14
  Defense Advanced Research Projects Agency (DARPA). *See* Defense Advanced Research Projects Agency (DARPA)
  denial-of-service attacks against, 18·3–18·5
  Directive 8570, 51·6
  FedRAMP, 71·46–71·47
  FISMA, 71·45
  Mil Standard 8500, 51·6
  NIST 800-53, 71·45–71·46
Department of Homeland Security (DHS):
  National Cyber Security Division, 38·15
  security planning and management procedures, 22·7–22·9, 22·26, 26·4
  Vulnerability Discovery ad Remediation Open Source Hardening Project, 51·18
Destruction, disclosure, use, and modification (DDUM), 3·18
Destruction of information, 3·2, 3·10–3·11, 3·14–3·15
Dial-up access to hardware, 4·14–4·15
Dial-up phone lines, 15·9
Dial-up servers, 25·5–25·6
Dialers, 16·8
Diffie-Hellman, 7·23–7·24, 7·36–7·37, 33·33, 37·16
Diffused IR, 6·12
Digital certificates, 7·31–7·35, 17·6, 17·9, 21·8
Digital coin payments, 21·11–21·12
Digital Equipment Corporation (DEC), 2·5
Digital Forensics Research Workshop (DFRWS), framework for ditigal investigation, 55·2–55·9
Digital Millennium Copyright Act (DMCA), 11·14–11·18, 11·22–11·23, 31·9, 35·6, 42·9–42·10, 42·12
Digital Music Access Technology (DMAT), 42·12
Digital rights. *See also* Intellectual property
  digital rights management (DRM), 26·15, 42·3, 42·12–42·17
  fundamental problems, 42·17–42·18
  overview, 42·1–42·2, 42·18–42·19
  piracy. *See* Piracy
  political opposition to, 42·15–42·17
  and privacy, 42·3
  privacy-enhancing technologies, 42·14–42·15
  published attacks, 42·17
  reverse engineering, 42·17
  tools for cracking, 42·17
Digital Signature Guidelines, 37·8–37·9

Digital Signature Standard (DSS), 37·23
Digital signatures:
Digital Signature Guidelines, 37·8–37·9
Digital Signature Standard (DSS), 37·23
and email, 21·7
and integrity checking, 46·4
and log management systems, 53·16
production programs, validating, 47·9
and protecting Web applications, 21·8
and public key cryptosystems, 37·4–37·5, 37·11, 37·13, 37·16, 37·19, 37·20, 53·16
source code, 38·8, 47·10
Digital subscriber line (DSL), 4·15, 5·4
Digital versatile disc, 4·9
Digital video discs (DVDs), 4·9, 36·1, 42·11
Digraphs, 7·12
Direct Access File System (DAFS), 57·4
Direct access storage devices (DASDs), 4·9–4·10
Direct attached storage (DAS), 36·4, 36·5
Direct memory access (DMA), 24·6
Direct sequence spread spectrum (DSSS), 6·12, 25·7
Directed beam IR, 6·12
Directory browsing, 21·18
Dirt and dust, 4·12
Disaster recovery:
backup/restore, 26·24–26·25
cold sites, 59·10
commercial recovery services, 59·12–59·13
and computer security incident response team, 56·7–56·8
data backup scenarios, 59·13
and gateway security devices, 26·24–26·25, 26·29–26·30
hot sites, 59·10, 59·16–59·18
implementation of plan, 59·20–59·21
internal redundancy, 59·11
mobile data centers, 59·11–59·12
overview, 59·1, 59·21
phases, 59·7–59·9
postincident analysis, 56·33–56·34
priority replacement agreements, 59·12
reciprocal agreements, 59·11
recovery strategies, 59·6–59·13
recovery tasks, 59·14–59·20
reserve systems, 59·13
security policies for, 66·11
scenarios, 59·3–59·6
testing of plan, 59·20–59·21
threats and threat assessment, 59·1–59·2. *See also* Threats
Disclosure of information, 3·2, 3·13, 3·16
Discretionary access control list (DACL), 24·18–24·19
Discretionary access controls, 9·2, 9·6, 9·9
Discs:
BD-R, 4·9
BD-RE, 4·9
compact disc read-only memory (CD-ROM), 4·9
holographic, 57·8–57·9
magnetic, 36·1
Discussion groups, 48·9–48·10
Disintermediation, 48·6
Disk space, log records for, 53·14–53·15
Distributed access control, 6·4
Distributed COM (DCOM), 21·13
Distributed denial-of-service (DDoS) attacks
defenses, 18·25–18·30
High Orbit Ion Cannon (HOIC), 18·22
history of, 18·15
HTTP Apache attack, 18·21
HTTP Unbearable Load King (HULK), 18·23
and information warfare, 14·21–14·22
and intrusion detection response, 27·11
Low Orbit Ion Cannon (LOIC), 18·22
and mobile code, 17·13
Mydoom, 18·22
overview, 18·2, 18·16–18·19
and real-time monitoring, 53·8–53·9
Shaft, 18·21
Stacheldraht, 18·15, 18·20
SubSeven, 18·22
terminology, 18·16–18·19
Tribe Flood Network (TFN), 18·15, 18·20
Tribe Flood Network 2 K (TFN2 K), 18·21
Trinoo (Trin00), 18·15, 18·19–18·20
Trinity, 18·15, 18·21–18·22
Distributed network traffic monitor, 18·29
Distributed polling, 6·14–6·15
DNSSEC, 5·26
Do not call list, 11·29
Document root, 21·17–21·18
Documentation:
changes, 39·17
computer security incidents, 56·17–56·21, 56·24
HIPAA, 39·16, 39·20
preservation of records, 48·36
and regulatory compliance, 39·20
software development, 38·13, 39·11, 39·18
Domain Name System (DNS):
attacks, 5·25–5·26, 18·10–18·11, 21·12
and block lists, 31·6
cache poisoning, 5·25–5·26
denial-of-service attacks (DoS), 5·13, 5·18–5·19
and host info (HINFO) resource record, 25·10
and network file system, 36·9
security, 5·26
server hierarchy, 5·26
and site names, 8·17
spoofing, 8·8, 17·11
updating DNS zones, 30·25
Domain names, and certification authority, 37·17–37·18

Dongles, 28·4, 28·13, 28·15–28·16, 42·5–42·6. *See also* Locks and door hardware; Smart cards
Dot dot attacks, 15·29–15·30
Double-entry bookkeeping, 1·3
Downloading software, 48·14
Downtime, 4·12–4·13
Driver's Privacy Protection Act, 69·14
Dropbox, 35·3–35·4
Due diligence:
  and code security, 38·5–38·5
  and information security framework, 3·14, 3·20
  and record keeping, 56·13
Dummy security devices, 23·8
Dumpster diving, 15·19, 19·6–19·7
Duties, separation of, 45·8–45·10, 47·4–47·5
D-Wave, 7·43
Dynamic Host Configuration Protocol (DHCP), 5·3, 5·26–5·27, 33·41, 34·7
Dynamic link libraries (DLLs), 47·10
Dynamic pricing, 48·24
Dynamic routing protocols, 5·27

E-cash, 21·11–21·12
E-commerce:
  applications design, 30·28
  auditing, 30·40–30·41
  business-to-business security services, 30·13–30·17
  business-to-customer security services, 30·3–30·4, 30·9–30·13
  counter-battery, 30·22
  ethical issues, 30·41–30·45
  and extranets, 32·16
  hold harmless, 30·22
  insurance policies, 60·15–60·17
  interruptions, 30·24–30·25
  and just-in-time production, 30·24
  and law enforcement cooperation, 30·20
  legal issues, 30·41–30·45
  loss of customers, 30·24
  operational requirements, 30·26–30·33
  overreacting to attacks, 30·21
  overview, 21·22, 30·2–30·3, 30·45
  PR image, 30·23–30·24
  responding to attacks, 30·21
  risk analysis, 30·22–30·25
  rules of engagement, 30·21
  security framework, use of, 30·9–30·18
  technical issues, 30·33–30·41
  threat and hazard assessment, 30·25
  vulnerabilities, 21·1–21·5, 21·8–21·22
  Web application system security, 21·5–21·8
  Website protection, 30·18–30·22
E-Sign Bill, 42·14
EAP RADIUS impersonation attack, 33·25
EAPoL (EAP Over LAN), 33·11–33·16
Earthlink, 20·18
Earthquakes and tsunamis, 22·15. *See also* Physical threats
Eavesdropping, 33·11, 33·21, 34·5–34·6
eBay, 48·25–48·26
EC-Council, 74·15–74·16
Echo, 4·6, 5·24, 18·13, 18·20
Echo messages, 18·8
Echo reply, 5·24, 18·20
Economic Espionage Act of 1996 (EEA), 11·7–11·8, 14·15–14·16
Edit checks, 47·16–47·17
EFF DES Cracker, 7·21, 7·38
Egress filtering, 18·14, 18·26
Electrical EPROMs (EEPROMs), 4·9
Electrical power. See Power failures and disturbances
Electromagnetic pulse attack (EMP), 22·20
Electromagnetic radiation (EMR), 25·5. *See also* TEMPEST (Transient ElectroMagnetic Pulse Emission Standard)
Electronic Communications Privacy Act (ECPA) of 1986, 11·30, 69·13–69·14
Electronic Data Interchange (EDI), 32·16
Electronic Frontier Foundation (EFF), 42·15–42·16
Electronic Media Management System, 42·13–42·14
Electronic Signatures in Global and National Commerce Act (E-Sign Bill), 42·14
Eliminate Vulnerable Code (eVc) Project, 51·18
Elliptic curve cryptography, 7·36–7·37
Email:
  anonymizing remailers, 48·4
  antivirus systems, 16·13
  appending services, 20·16–20·17
  archiving, 48·36
  authentication, 20·25–20·26
  autoforwarding, 48·22–48·23
  black holes, 20·2–20·21
  blacklists, 20·4
  block lists, 20·20–20·21
  CC and BCC functions, 48·20–48·21
  centralized distribution lists, 48·18–48·19
  chain letters, 48·10–48·11, 48·15
  changing email addresses, 48·30
  corporate identifiers, 48·4–48·5
  data leakage, 13·9
  digital signatures, 21·7
  disclaimers, 48·17–48·18
  distribution lists, 48·20–48·21
  employee misuse of, 48·3–48·4
  first impressions, 48·17
  flaming, 48·3–48·4
  forwarding, 48·15–48·16
  headers, 20·4–20·6, 48·9
  HTML, 17·3, 20·15, 28·2, 48·19–48·20
  illegal activities, 48·3
  inefficient use of, 48·15–48·22
  and Internet hoaxes, 48·6–48·12, 48·44

list servers, 48·23
mail-bombing, 18·6–18·8, 48·3, 48·43
mail storms, 18·7, 48·22–48·24, 48·43–48·44
Mail Transfer Agent (MTA), 20·3, 20·13
mass mailing, 20·14–20·16
monitoring tools, 19·19
multilevel marketing schemes, 48·11
and pedophiles, 48·13
permission issues, 20·16–20·17
pharming, 46·11
phishing, 46·11
Ponzi schemes, 48·10–48·11
private e-mail in the workplace, 48·21–48·22
reply all, 48·21
Simple Mail Transfer Protocol (SMTP), 5·28, 17·15, 20·3–20·4, 20·25
spam. *See* Spam
subject line, 48·16
and threats of physical harm, 48·12–48·13, 61·5–61·6
transfer standards, 5·28
Unified Threat Management (UTM), 31·9
worms, 16·8
Email storm 18·7
Emergency Management Assessment Program (EMAP), 22·9, 23·3
Emergency power, 23·30–23·36
Emergency Services Working Group, 6·19
Emerging technologies, 30·41
Employees:
abuse, opportunities for, 45·3
access issues, 45·4
background checks, 16·12, 45·2–45·3
behavioral changes, 45·7–45·10
blogs, 48·5
capability and responsibility matrices, 45·10–45·11
career advancement, 45·6
and computer crime, 45·1–45·2
confidentiality agreements, 45·16
confirmation bias in, 19·12
cross-training, 45·4–45·6
employment agreements, 11·4, 45·3
extranet services for, 32·17
fiduciary duties, 11·4, 11·30
firings, 45·12–45·16
health and safety issues, 22·3, 23·42–23·43
identification and authentication issues, 28·17
indispensable, 45·4–45·6
Internet use, policies on, 48·45. *See also* Internet
malicious use of compromised username/password-based authentication credentials, 33·42
and management practices, 45·3–45·12, 45·16
motivation, 63·12–63·14
noncompetition agreements, 45·16
overview, 45·16–17
personal Websites, 48·5
productivity, threats to, 48·15–48·30
resignations, 45·12
separation of duties, 45·8–45·10
shiftwork, 56·31–56·32
and social engineering. *See* Social engineering
and social networking, 48·5
social psychology, use of in implementing security policies. *See* Social psychology
supervision of, 63·14–63·15
termination, 45·12–45·16
and unauthorized security probes, 45·11–45·12
vacation time, 45·6–45·7
Web monitoring, 31·3–31·4
Employment agreements, 11·4, 45·3
Enclaves, 17·2, 17·14
Encrypting File System (EFS), 25·11
Encryption:
additional decryption keys (ADKs), 36·12
Advanced Encryption Standard (AES), 7·39, 7·43–7·44
algorithms, 7·2, 37·16, 37·23
and applications design, 30·28
authenticity and trust, 7·26–7·27
brute force cryptanalysis, 7·9–7·11, 7·17, 21·12, 37·22
ciphertext, 7·2–7·3, 37·2–37·3
data backups, 36·6
data encryption, 37·23–37·25
data storage, 36·10–36·15
database, 36·13–36·15
biometric encryption, 29·19, 29·21
brute force cryptanalysis, 7·9–7·11, 21·12
Caesar cipher, 7·8, 7·7–7·11
cell phones, 15·12
ciphers, 4·17, 7·2, 7·6–7·16
ciphertext, 7·2–7·3, 7·9, 37·2–37·3
codes, 7·2
communications, 7·27–7·35
cryptography, 4·16–4·17, 7·3–7·16, 37·2
cryptology, 7·2
cryptoviruses, 16·7
Data Encryption Algorithm (DEA), 7·20, 7·38
Data Encryption Standard (DES), 7·2, 7·16, 7·17–7·22, 25·11
decryption, 7·2–7·3, 7·9–7·10, 7·26, 7·37
defined, 31·2
digital certificates, 7·31–7·35
e-commerce security services, 30·6
elliptic curve cryptography, 7·36–7·37
Encrypting File System (EFS), 25·11
of instant message communications, 35·12–35·13
ISO/IEC/ITU 9594-8 (X.509). *See* X.509 certificate format
key-exchange problem, 7·8, 7·23
key escrow, 36·11
keys, 4·17, 7·2
keyspace, 7·3, 7·10, 7·21, 15·16–15·17

Encryption (*Continued*)
log records, 53·16
Message Authentication Code (MAC), 7·4, 7·30, 38·8
mobile data systems, 1·19
monoalphabetical substitution ciphers, 7·11–7·12
and network security, 26·13–26·14
one-time pad, 7·17–7·18
overview, 7·1–7·2
and password cracking, 15·25
passwords, 28·9–28·11
plaintext, 7·2–7·3, 7·9
Point-to-Point Encryption, 25·8
private keys, 7·5, 7·8
public key, 5·28, 7·5, 7·21–7·27, 28·10–28·11
Public Key Infrastructure (PKI). *See* Public Key Infrastructure (PKI)
quantum cryptography, 7·39–7·44
RSA algorithm, 7·25–7·27, 7·36–7·38
Secure Sockets Layer (SSL), 7·28–7·30
and security controls, 3·14
SEEK, 7·27
smart phone encryption, 36·15
steganography, 7·44–7·45, 31·12
substitution ciphers, 7·8
symmetrical encryption, 7·8
terminology, 7·2–7·3, 7·8
Transport Layer Security (TLS), 7·28–7·30
transposition, 7·18–7·20
trust, 37·2
and virtual private networks (VPNs), 32·11
and VoIP, 34·9–34·10
Web applications, protecting, 21·7
and Web monitoring, 31·11–31·12
wireless networks, 15·12, 25·7–25·8, 33·8–33·9
XOR (Exclusive-Or), 7·15–7·16
End user license agreement (EULA), 16·9
Endangerment, 3·2, 3·14, 3·17–3·18
Endpoint attacks, 33·24
Endpoint supplicant configurations, 33·36–33·37
Energy anomalies as security hazard, 3·18
Energy systems and information warfare, 14·9–14·11
Engressia, Josef ("Joybubbles"), 2·8
Enhanced 911, 34·2–34·3
Enhanced Interior Gateway Routing Protocol (EIGRP), 5·27
Enterprise JavaBeans (EJB), 21·13
Enterprise patch management technologies, 40·7–40·10
Enterprise resource planning (ERP) systems, 32·16
Enumerating, 55·13–55·14
Environmental alarms, 23·18
Environmental threats, monitoring, 53·18
Equipment cabinets, 23·10
Equity Funding fraud, 2·9–2·10
Errors:
boundary condition violations, 38·11
calculation errors, 38·11, 39·9
codes and coding, 16·13, 21·7, 38·10–38·15, 39·8–39·12
control (command) structure errors, 38·14, 39·11–39·12
biometric authentication crossover error rate (equal error rate), 29·16
data input errors, 52·14
functionality errors, 38·13, 39·11
initialization error, 38·11, 39·8–39·9
interapplication conflicts, 38·12, 39·10
load condition errors, 38·12, 39·9
logic flow errors, 38·11, 39·9
output format errors, 38·15, 39·12
parameter passing errors, 38·11–38·12, 39·9
performance errors, 38·14–38·15, 39·12
performance (speed) errors, 38·14–38·15
program conflict errors, 38·
race conditions, 38·12, 39·9, 39·14
resource exhaustion, 38·12
software errors, types of, 38·10–38·15, 39·8–39·12
E.T. applications, 48·14
Ethereal, 33·45
Ethernet, 1·10, 5·13, 6·3
Ethernet hubs, 6·25
Ethernet II, 6·21–6·22
Ethernet Working Group, 6·17
Ethics:
and anonymity, 70·24–70·26
and computer crime, 12·14–12·16
customer monitoring, 30·42–30·43
defined, 43·1–43·2
and dummy security devices, 23·11
importance of, 43·1–43·2, 43·7–43·8
and penetration testing, 46·8
principles, 43·2–43·7
and surveillance systems, 23·20
and Website management, 30·42–30·43
European Network and Information Security Agency (ENISA), 49·32
European Union Data Protection Directive 95/46/EC, 69·8–69·11
implementation of, 69·10–69·11
international data-transfer restrictions, 69·9–69·10
requirements, 69·8–69·9
U.S./EU safe harbor, 69·22–69·23
European Union Data Protection Supervisor (EDPS), 69·11–69·12
European Union Telecommunications Directive, 69·11
Evanescent media, 42·10
Event logging, 19·18
Events, 8·4–8·11. *See also* Security incidents

Evidence:
chain of custody, 55·12, 55·17
collecting, 55·10, 55·16
corroboration of, 55·12, 55·17
Exception reports, 53·24
Executive Order—Improving Critical Infrastructure Cybersecurity, 60·20
Exhaustive search, 15·16
Expert witnesses:
admissibility of scientific evidence, 73·1–73·4
appearance, 73·6
background, 73·1–73·2
Daubert v. Merrell Dow Pharmaceuticals, 73·2–73·4
Federal Rules of Evidence, 73·2
fees, 73·5
General Electric Co. v. Joiner, 73·3
Kumho Tire Co. v. Carmichael, 73·3
overview, 73·6
preparation for testimony, 73·4–73·6
qualifying as, 73·4–73·5
state law, 73·4
written report, 73·5–73·6
Exploiting curiosity or naïveté, 19·8
Exposure effect, 19·12–19·13
Extended Key Usage, 37·6
Extensible Authentication Protocol (EAP), 33·4, 33·7, 33·11–16
Extensible Markup Language (XML), 17·3, 21·13
External interrupts, 4·7
External service zones, 26·18–26·19
Extortion, 2·12
Extranets, 26·19, 32·15–32·19

414s, 2·23
Fabric Login (FLOGI), 36·8
Facebook, 48·5, 69·27
Facial recognition scanning, 23·15
Fair and Accurate Credit Transaction Act of 2003, 11·29
Family Education Rights and Privacy Act (FERPA), 67·3–67·4
Faraday cages, 22·20, 25·5, 53·5
Fault injection, 38·18
Federal Bureau of Investigation (FBI):
adversarial matrix of behavioral characteristics, 55·18–55·20
adversarial matrix of operational characteristics, 55·20–55·21
adversarial matrix of resource characteristics, 55·20, 55·22
Bot Roast II, 17·13
InfraGard, 1·18–1·19
Project Megiddo report, 22·5
reporting threats and incidents to, 61·7
and use of spyware, 17·4
Federal Communications Commission (FCC), 34·2
Federal Emergency Management Agency (FEMA):
guidelines for security management, 23·5
Publication 386-1, *Getting Started,* 23·43
Publication 386-2, *Understanding Your Risks,* 23·43
Publication 386-3, *Developing the Mitigation Plan,* 23·43
Publication 386-4, *Bringing the Plan to Life,* 23·43
Publication 386-5, *Using Benefit-Cost Review in Mitigation Planning,* 23·44
Publication 386-6, *Integrating Historic Property and Cultural Resource Considerations into Hazard Mitigation Planning,* 23·44
Publication 386-7, *Integrating Manmade Hazards into Mitigation Planning,* 23·44
Publication 386-8, *Multi-Jurisdictional Mitigation Planning,* 23·44
Publication 386-9, *Using the Hazard Mitigation Plan to Prepare Successful Mitigation Projects,* 23·44
security planning and management procedures, 22·7–22·9
State and Local Mitigation Planning guides, 23·43
Federal Financial Institutions Examination Council (FFIEC), 67·3
Federal Information Processing Standards (FIPS), 51·6, 51·23, 54·9, 67·9–67·12
Federal Information Security Management Act (FISMA), 54·4, 54·5–7, 71·45
Federal Rules of Civil Procedure, 67·5
Federal Rules of Evidence, 73·2
Federal statutes under consideration in 2013, 60·22–60·23
Federal Trade Commission (FTC):
FTC Act Section 5, 69·29
and Gramm-Leach-Bliley Act enforcement, 64·7, 69·17
and healthcare information breaches, 71·24
reporting identity theft to, 61·6
Federated Identity Management (FIdM), 32·19
FedRAMP program, 51·24–51·26, 71·46–71·47
Feistel ciphers, 7·20
Fiber Distributed Data Interface (FDDI), 6·14
Fiber optic cable, 15·11
Fiber optic perimeter alarms, 23·17
Fiber Optics Technology Advisory Group (FOTAG), 6·18
Fibre channels, 36·4, 36·7–36·8
File close log record, 53·13
File infector viruses, 16·4
File input-output log record, 53·13–53·14
File open log record, 53·13
File sharing, 11·23–11·24, 24·11–24·12
File system logging, 53·20
File Transfer Protocol (FTP), 5·28, 21·12, 30·39

Filtered queries, 53·21
Financial Institution Reform Recovery and Enforcement Act (FIRREA), 64·8
Fingerprint scanners, 23·14–23·15
Fire and smoke, 22·20, 23·38–23·41
Firewalls:
  access control lists, 26·7
  appliances, 26·7–26·8
  application layer gateways (ALGs), 26·6
  changing architectures, 26·4–26·5
  deployment, 26·18–26·20
  gateway security device as replacement for, 26·25
  and host environment, 26·6–26·7
  host-based, 26·9
  and internal partitions, 30·27
  intrusion detection. *See* Intrusions
  and IP addresses, 5·15
  and malicious code, 16·13
  and mobile code, 17·4
  monolithic, 30·35
  next-generation firewall (NGFW), 26·4, 26·6
  operating system security, 21·21
  overview, 26·2–26·4
  packet filtering, 26·5
  platforms, 26·7
  as protection against denial-of-service attacks, 18·26–18·27
  routers, 5·3, 26·7
  stateful inspection, 26·5
  unified threat management (UTM), 26·4, 26·6
  usability, 26·30
  and use of RFC 1918 addresses, 30·34
  and virtual private networks, 32·6
  and VoIP, 34·7
  and Web application protection, 21·7–21·8, 26·4
  and Web monitoring, 31·9–31·10
  Windows XP, 25·11
Firmware updates, 40·5
First Amendment, U.S. Constitution, 48·4, 72·12–72·21
Flag fields, 5·20
Flash, 26·16
Flash drives, 1·19, 36·1. *See also* Thumb drives
Flash memory, 4·9
Flexible Authentiaction via Secure Tunneling (FAST), 33·12
Flooding, 22·15. *See also* Physical threats
Flu-shot hoax, 2·12
Flux bot, 15·32
Focus groups and computer crime research methods, 10·9–10·10
Footprinting, 55·13
Foreign Intelligence Surveillance Act (FISA), 69·15–69·16
Forum of Incident Response and Security Teams (FIRST), 56·36
Frame Relay networks, 5·5
Frames, 5·8, 5·12–5·14
Fraud:
  advance-fee fraud, 2·21, 16·13
  digital certificates, 17·9
  and get-rich-quick schemes, 48·11–48·12, 48·44–48·45
  loss of possession of information 3·14
  Nigerian 419 fraud (advance-fee fraud), 2·21, 16·13, 19·9
  and online auctions, 48·25–48·26, 48·41
  reporting, 61·6
Free Software Foundation, 42·16
Freeware, 48·42
Frequency analysis, 7·11–7·13
Frequency hopping spread spectrum (FHSS), 6·12, 15·12, 25·7
Friendster, 48·5
  Functionality errors, 38·13, 39·11
Fundamental attribution error, 19·13
Fuzzing, 38·18–38·19, 39·16

Gases as security hazard, 3·17
Gates, Bill, 1·9
Gateways. *See also* Routers
  application layer gateways (ALGs), 26·6
  Border Gateway Protocol (BGP), 5·23, 5·27
  Common Gateway Interface (CGI), 15·27–15·28, 21·2–21·4, 21·13–21·15, 21·17
  Enchanced Interior Gateway Routing Protocol (EIGRP), 5·27
  security devices. 26·1–26·34. *See also* Firewalls
  terminology, 5·7
General Agreement on Tariffs and Trade (GATT), 11·35
General Electric Co. v. Joiner, 73·3
General Public License (GPL), 11·33–11·34
Genetic Information Nondiscrimination Act, 69·27
Germany, 44·9, 51·6
Glass box testing approach, 39·15
Global Information Assurance Certification (GIAC), 74·16–74·18
Global Positioning System (GPS), 69·26
Globalization, 69·1–69·2. *See also* Outsourcing
Google:
  background checks, 45·3
  and collaboration tools, 35·3–35·4, 35·18. *See also* Collaboration tools
  Perfect 10 litigation, 11·23
  report on malware, 15·31
  Safe Search, 31·13
Google Docs, 35·3
Google Glass, 69·26
Government:
  biometric authentication, 29·22
  electronic communications and privacy, 11·30

and information warfare, 14·8–14·9
and VoIP compliance, 34·2–34·3
Gramm-Leach-Bliley Act of 1999:
applicability, 64·6–64·7
compliance, 64·8–64·9
consumer defined, 64·8
customer defined, 64·8
and data classification, 67·4–67·5
and data protection, 26·3
and documentation of changes, 39·20
enforcement, 64·7–64·8, 64·10, 64·14
evaluation procedures, 64·11–64·13
examination procedures, 64·14
and healthcare information protection, 71·15, 71·35–71·39
nonpublic personal information, 60·21–60·22
overview, 64·6
privacy notices, 64·9–64·10
provisions of, 54·5–54·6, 64·9, 69·17–69·18
Safeguards Rule, 64·9–64·10
and security audits, 54·2, 54·5
and security planning, 23·3
and VoIP compliance, 34·2
Gray box testing, 38·18
Green Card Lottery spam, 2·20–2·21
Groupthink, 50·20–50·21
Grover's algorithm, 7·41–7·42

Hackers:
attackers, categories of, 8·16
and computer criminals, 12·2–12·4, 12·16–12·17
and e-commerce systems, 21·1–21·5. See also E-commerce
ethics, 12·14–12·16
and link analysis, 55·22–55·23. *See also* Cyberinvestigation
motivation, 4·21
penetration techniques, 15·3–15·4, 15·6
and social engineering, 19·3. *See also* Social engineering
and software development, 39·20
support groups, 15·39
workplace issues, 48·33, 48·45
Hackers underground:
Anonymous, 2·28
Chaos Computer Club (CCC), 2·23
Cult of the Dead Cow (cDc), 2·23–2·24
the 414s, 2·23
Legion of Doom, 2·24–2·25
L0pht Heavy Industries, 2·27
Masters of Deception, 2·25–2·26
Operation Sundevil, 2·26
Phrack, 2·25
Russian Business Network (RBN), 2·27–2·28
Steve Jackson Games, 2·26–2·27
2600: The Hacker Quarterly, 2·24
Unlimited Operations, 2·28–2·29
Haephrati Trojan, 2·14
Handbook for Computer Security Incident Response Teams (CSIRTs), 56·5
Handshake protocols, 7·28–7·29, 25·10
Hardware. *See also* Computers
antipiracy techniques, 42·5–42·12
backup, 4·17–4·12
binary design, 4·2–4·4
cryptography, 4·16–4·17
data communications, 4·13–4·16, 4·23–4·24
interrupts, 4·7–4·8
memory and storage data, 4·8–4·10
parity, 4·4–4·5
personal computers, 4·20–4·25
physical and environmental threats, 4·11–4·13
recovery procedures, 4·20
security checklist, 4·25–4·27
threats, 24·3
time functions, 4·10
tokens, 28·15–28·16
Hardware Trojans, 2·14–2·15
Harpooning, 17·4
Harrison, Ruzzo, and Ullman (HRU) Result, 9·5
Harvesting, 20·7–20·8
Hash totals, 46·4
Hate groups, 48·34, 48·37
Hazardous materials, 22·20
Health Information Technology for Economic and Clinical Health (HITECH), 60·21, 71·16–71·17, 71·28, 71·30–71·31, 71·41–71·45
Health Information Trust Alliance (HITRUST), 71·48–71·50
Health Insurance Portability and Accountability Act (HIPAA). *See also* Healthcare security and privacy
breach notification enhancements, 71·29–71·30
Breach Notification Interim Rule, 71·24–71·27
comparison of federal and state laws, 71·34
compliance, 71·31–71·34
electronic communications and privacy, 11·29
and data classification, 67·4
documentation, 39·16, 39·20
enforcement enhancements, 71·30–71·31
privacy and confidentiality standards, 69·20–69·21, 71·15–71·40
privacy enhancements, 71·29
Privacy Rule, 71·19–71·21
Privacy, Security, and Breach Notification Rules, 71·27–71·28
privacy and security principles, 71·18–71·19
protected health information, 60·21
related federal guidelines, 71·34
and risk management, 26·3
and security audits, 54·2, 54·7
security enhancements, 71·28–71·29
and security planning, 23·3
Security Rule, 71·21–71·24, 71·28
and VoIP compliance, 34·2–34·3

Health threats, 22·3, 22·16, 23·42–23·43. *See also* Physical threats
Healthcare accrediting commissions, 71·33
Healthcare security and privacy:
  attack vectors, 71·12–71·13
  big data trends, 71·53–71·54
  breaches, 71·5–71·12
  evolving and emerging government initiatives, 71·47–71·48
  federal laws, 71·14–71·16
  FedRAMP, 71·46
  FISMA, 71·45
  HIPAA information protection regulations, 71·17–71·40
  HITECH Act, 71·41–71·45
  personalized medical technology, 71·53
  private sector initiatives, 71·48–71·51
  mobile and personal technology, 71·52
  motivation for, 71·4
  NIST 800-53, 71·45–71·46
  overview, 71·2–71·3, 71·55–71·56
  technology trends, 71·51–71·55
  threats and vulnerabilities, 71·4–71·14
  virtualized technology, 71·53
Help files, 44·14
Helpdesk, 19·4–19·5, 56·21, 56·27, 56·29
Heuristic malicious code detection, 16·8–16·9
Heuristics, 19·12, 41·8–41·9
Hidden fields, 21·16–21·17
High-energy radio-frequency (HERF) weapons, 22·19
High-Layer LAN Protocols Working Group, 6·17
High Level Interface (HILI), 6·17
High-Tech Crime Network (HTCN), 74·18
Honeypots, 13·10–13·11, 23·11
Host:
  defined, 5·2
  rogue or counterfeit, 36·10
  trusted, 36·9
Host Intrusion Prevention System (HIPS), 26·18, 53·18
Host Identity Payload (HIP), 18·30
Hostile work environment, 48·33–48·36
Hotspot attacks, 33·27–33·28
HTML. *See* Hypertext Markup Language (HTML)
HTTP. *See* Hypertext Transfer Protocol (HTTP)
HTTPS, 30·29, 30·30, 30·35, 30·41, 32·16
Hubs, 6·25
Human–machine interface (HMI), 53·7–53·8
Humidity, 4·11–4·12, 23·37
Hurricane Katrina, 22·2, 22·9, 22·15
Hyperlinks, 11·23, 44·13–44·14
Hypertext Markup Language (HTML), 5·28, 15·28, 17·3, 28·2, 44·13–44·14, 48·19–48·20
Hypertext Processing (CGI/PHP), 21·13
Hypertext Processor (PHP), 15·30
Hypertext Transfer Protocol (HTTP), 5·10, 5·28, 21·7, 21·10, 21·12, 21·16–21·18, 30·12, 30·15, 30·30, 30·35
Hypertext Transfer Protocol Daemon (HTTPD), 15·28
Hypothesis testing, 10·5

IBM:
  Data Encryption Standard (DES), 7·2, 7·16, 7·17–7·22
  denial-of-service (DoS) attack, 18·3
  Electronic Media Management System, 42·13–42·14
  Lucifer product cipher, 7·19–7·21
  memory protection, IBM System/370, 24·6
  Millipede, 57·12
  Model 650, 1·6
  Model 5100 PC, 1·9
  Predictive Failure Analysis (PFA), 4·10
  quantum computing, 7·41–7·42
  System/370, 24·6
  Token Ring, 6·9, 6·14
  Virtual Machine technology, 17·14
  Virtual Tape Server, 57·10
ICSA Labs, 51·21
Identification:
  defined, 28·2
  e-commerce security services, 30·6
  federal employees and contractors, 28·17
  issues, 28·17
  and operations security, 47·6
  software versions, tracking, 47·7–47·9
  verification, 29·5–29·7
Identity theft, 2·7–2·8, 70·2
IEEE 488 standard, 44·3
IEEE 802 standards, 4·16, 5·3, 5·13, 6·14, 6·17–6·24, 25·7, 33·4–33·7, 33·9–33·20, 33·23–33·24, 33·46–33·49
Impersonation, 2·5–2·8, 19·4–19·5
Income, loss of, 60·12–60·14
Incremental information leveraging, 15·6–15·7
Industrial control systems (ICSs), 26·19, 53·26–53·27
Industrial espionage, 2·29–2·31
Inference, in statistics, 10·4, 10·8
Information assurance (IA):
  accreditation, 74·8–74·9
  BYOD policy, 75·3–75·4
  certification, 74·3–74·4, 74·6–74·8, 74·10–74·23
  changing posture, 75·2–76·4
  complexity, 75·1–75·2
  computers at risk, 1·14–1·18
  cybersecurity professionals, 74·1–74·8
  decentralized security, 75·8–75·9
  demilitarized zone (DMZ) model, 75·3
  government's role, 1·13–1·19
  healthcare security, 75·5
  honeycombs, 75·3

IA standards, 1·13–1·14
InfraGard, 1·18–1·19
legal and jurisdictional issues, 75·5–75·7
overview, 74·24, 75·1
preparing for exams, 74·23
quantum computing, 75·4–75·5
quarantine, 75·3
recognition by government agencies, 74·2
and recoverability of data, 36·10–36·11
ubiquitous technologies, 75·7–75·8
workforce training, 74·5–74·6
Information infrastructure:
access control. *See* Access control
defensive strategies, 23·4
facilities design, 23·21–23·41
overview, 23·2
physical site security. *See* Physical site security
protecting, 23·5–23·9
responsibility for, 23·5
strategic planning, 23·3–23·5
targets, 23·4
threats to, 22·2, 22·3. *See also* Physical threats
Information life cycle management (ILM), 67·2
Information security (IS):
cost-benefit analysis, 23·44–23·45
implementation, accountability, and follow-up, 23·46
mitigation plan, 23·43–23·44
net present value of, 63·5
security response plan, 23·45–23·46
and strategic goals, 63·4–63·5
Information security (IS), new framework proposal:
acts that cause loss, 3·2–3·3, 3·10–3·19
components of, 3·2–3·3
need for, 3·1–3·2
objectives of information security, 3·3
purpose of, 3·20, 3·23
safeguard functions, 3·3, 3·19–3·20
safeguard selection methods, 3·3, 3·20
security elements, 3·2, 3·3–3·9
sources of loss, 3·2, 3·10–3·19
terminology, 3·9–3·10
threats, assets, vulnerabilities model, 3·2, 3·20–3·22
Information security administrators, 47·4–47·5
Information Security Breaches Survey (U.K.), 13·3
Information Security Policies Made Easy (ISPME), 44·9
Information Security Policies and Procedures, 44·10
Information-sharing initiative—implementation plan, 61·15–61·17
Information systems (IS):
audits, 35·8, 35·21
evolution of systems, 1·3–1·13
family of computers, 1·7–1·8
first "worm," 1·9
history of, 1·3–1·13
information assurance (IA), government recognition of, 1·13–1·19
infrastructure. *See* Information infrastructure
interconnection, 1·11–1·12
Internet and the World Wide Web, 1·12–1·13
large-scale computers, 1·4–1·5
local area networks (LANs), 1·10–1·11
location of equipment, 61·15–61·17
medium-size computers, 1·5–1·6
microprocessors, 1·8
network, first, 1·8
overview, 1·1–1·3
personal computers (PCs), first, 1·8–1·9
productivity enhancements, 1·9
punched-card system, 1·4
real-time, online systems, 1·7
recent developments, 1·19
security, overview, 1·1, 15·1–15·2
security breach prevention and mitigation, 35·5–35·7
small-scale computers, 1·6–1·7
supervisory control and data acquisition (SCADA), 1·13, 14·9
telecommuting, 1·12
time sharing, 1·7
transistors and core memory, 1·7
virtualization and the cloud, 1·13
Information Systems Audit and Control Association (ISACA), 54·4
Information Technology Management Reform Act of 1996 (Clinger-Cohen Act), 54·6
Information Technology Security Evaluation Criteria (ITSEC), 51·10
Information theft, 15·35–15·36
Information warfare (IW):
biological and chemical weapons, 14·25
and China, 14·17–14·19, 16·3
and commerce, 14·12–14–13
and commercial off-the-shelf software, 14·3
and communications, 14·16
cooperative efforts, 14·28–14·29
corporations as victim of, 14·21
and critical infrastructure, 14·2–14·3
cryptography, 14·23
cyberterrorists, 14·19–14·21
defined, 14·1
denial of service (DoS) and distributed denial of service (DDoS), 14·21–14·22
destabilization of economic infrastructure, 14·16–14·17
and energy systems, 14·9–14·11
financial disruptions, 14·12–14·14
forceful defenses, 14·26–14·27
goals and objectives, 14·5–14·17
and hardware Trojans, 2·14–2·15
in-kind counterattacks, 14·27
integration of cyberwarfare into military planning, 14·27–14·28

Information warfare (IW) (*Continued*)
international and corporate espionage, 14·15–14·16
and law enforcement, 14·15
legal defenses, 14·25–14·26
malicious code, 14·22
and medical security, 14·14–14·15
overview, 14·2
physical attacks, 14·24–14·25
psychological operations (PSYOP), 14·23–14·24
rebuttal, 14·4–14·5
technical defenses, 14·27
and transportation systems, 14·11–12
views on, 14·3–14·4
weapons inadvertently provided, 14·25
weapons of mass destruction, 14·25
InfraGard, 1·18–1·19, 61·12, 61·15
Infrared (IR), 6·12
Infrastructure:
attacks, 14·5
contract, 69·29–69·30
destabilization of economic infrastructure, 14·16–14·17
extranet, 32·17
information infrastructure. *See* Information infrastructure
and information warfare (IW), 14·2–14·3, 14·9–14·11
local area network (LAN) security, 25·3–25·9
maintenance and repair, 23·9–23·10
virtual private network (VPN), 32·6, 32·12–32·13
Ingress filtering, 18·14
Initialization errors, 38·11, 39·8–39·9
Injection laser diode (ILD), 6·10
Input/Output (I/O), 4·7, 24·2, 24·4–24·6
Insiders, information technology:
and classification of computer criminals, 12·19
mitigating threats, 13·8–13·11
overview, 13·1–13·2
threats, 13·2–13·8
Installation, 23·9–23·10
Instant messaging (IM):
business threats, 35·9–35·10
loss of information, 35·10
and need for security, 35·1
overview, 35·2, 35·9, 35·22–35·23
safe messaging, 35·13–35·14
security incident response, 35·13
and social engineering, 19·9
Institute for Credentialing Excellence, 74·9
Institute for Electronics and Electrical Engineers (IEEE), 6·17
IEEE 802 standards. *See* IEEE 802 standards
IEEE 488 standard, 44·3
IEEE 802 standards. *See* IEEE 802 standards
Institute of Internal Auditors (IIA), 74·19
Instrument validation, 10·10
Insurance:
brand image, damages to, 60·23–60·26
business interruption, 60·17
claims made coverage, 60·8–60·9
commercial general liability (CGL) policies, 60·3–60·9
and compliance with standards, 22·7
consumer issues, 60·26–60·27
costs for recovery, 60·14–60·15
covered claims, 60·10
crime/fidelity coverage, 60·15
damage from employees, 60·14
direct physical loss, 60·12
duty to defend, 60·9–60·10
e-commerce policies, 60·15–60·17
general liability issues, 60·3–60·4
exclusions, 60·11
extensions of coverage, 60·11
first-party coverage, 60·11–60·12
historical background, 60·2
identity theft, 60·17–60·23
indemnity, 60·9–60·10
intellectual property coverage, 60·3–60·12
loss of income, 60·12–60·14
need for, 60·3
occurrence coverage, 60·8–60·9
principal, 60·10
prior acts coverage, 60·11
privacy breaches, 60·17–60·23
property coverage, 60·12–60·15
Intangible assets:
insurance coverage for loss or damage, 60·7–60·8
intellectual property. *See* Intellectual property
Integrated Services LAN (ISLAN) Working Group, 6·18
Integrity:
checking, 46·4
data, 24·4
e-commerce security services, 30·7
and encryption, 7·4. See also Encryption
firewalls and gateway security devices, 26·20
as source of loss, 3·2, 3·5–3·6, 3·8–3·10, 3·12
Intellectual property:
Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS), 11·35–11·39
America Invents Act of 2011 (AIA), 11·39
circumvention of technical measures to secure copyrights, 11·15–11·18
Computer Fraud and Abuse Act (CFAA), 11·27–11·33
and contracts for protecting technology, 11·3–11·5, 11·25–11·26
Digital Millennium Copyright Act, 11·14–11·18, 11·22–11·23, 31·9
Electronic Communications Privacy Act (ECPA), 11·30

insurance coverage, 60·3–60·12
and open source, 11·33–11·34
outsourcing risks, 68·17–68·23
overview, 11·2–11·3, 11·39
patent trolls, 11·43–11·44
piracy, 11·20–11·24
trademarks, 42·2
Preventing Real Online Threats to Economic Creativity and Theft of Intellectual Property (PROTECT IP) Act (PIPA), 40–41
Stop Online Piracy Act (SOPA), 11·41–11·43
Stored Communications Act, 11·32–11·33
terms of use, 11·25–11·26
trade secrets. *See* Trade secrets
trespass, 11·25
unauthorized intrusions, legal remedies, 11·24–11·33
Wiretap Act, 11·30–11·32
World Intellectual Property Organization Copyright Treaty, 11·15
Interapplication conflicts, 38·12
Internal honeypots, 13·10–13·11
Internal service zones, 26·19
International Association of Computer Investigative Specialists (IACIS), 74·18–74·19
International espionage, 14·15–14·16
International Information Systems Security Certification Consortium (ISC)^2, 74·20–74·22
International law:
intellectual property, 11·35–11·39
and VoIP, 34·3
International Organization for Standardization (ISO):
British Standard 7799, 54·11–54·12
Common Criteria (CC), 51·3–51·4, 51·26–51·28
ISO 13485, 71·51
ISO 14971, 71·51
ISO 17799, 38·20, 44·4, 54·2
ISO 27000 Series, 54·4, 54·11–54·12
ISO 27001, 54·2
ISO 27002, 54·2
ISO 27799:2008, 54·13
ISO 9000, 51·7, 51·11
ISO/IEC 15408, 38·20
ISO/IEC 17011, 74·8
ISO/IEC 17799:2005, 44·3, 54·12
ISO/IEC 21827, 51·11
ISO/IEC 27000 series, 54·11–54·14
ISO/IEC 27000:2009, 54·12
ISO/IEC 27001:2005, 54·12
ISO/IEC 27002:2005, 54·12
ISO/IEC 27003:2010, 54·13
ISO/IEC 27003:2011, 54·13
ISO/IEC 27004:2009, 54·13
ISO/IEC 27005:2011, 54·13
ISO/IEC 27006:2011, 54·13
ISO/IEC 27007:2011, 54·13
ISO/IEC 27011:2008, 54·13
ISO/IEC 27034-1, 51·8, 51·13
ISO/IEC FDIS 27013, 54·13
ISO/IEC FDIS 27014, 54·13
ISO/IEC TR 27008:2011, 54·13
ISO/IEC TR 27015, 54·13
ISO/IEC WD 15443, 38·20
ISO/IEC WD 27016, 54·13
and OSI, 5·11
overview, 54·10–54·11
International Telecommunications Union–Telecommunications Standards Sector (ITU-T), 5·11
Internet:
addiction, 48·15, 48·28, 48·38
disconnection, 30·40
and dissemination of incorrect information, 48·5–48·6
fraud, reporting, 61·5–61·6
games, 48·29–48·30, 48·40
get-rich-quick schemes, 48·11–48·12, 48·15, 48·44–48·45
history, 1·12–1·13
hoaxes, 48·6–48·12, 48·44
online auctions, 48·15, 48·25–48·26, 48·41
online dating, 48·28–48·29, 48·39–48·40
online gambling, 48·15, 48·26–48·27, 48·41–48·42
online monitoring technology, 69·25–69·26
online shopping, 48·15, 48·24–48·26, 48·40–48·41
overview, 5·7–5·9, 48·2
and pedophiles, 48·13
and reputation damage, 48·2–48·12
system and network penetration through Websites, 15·27–15·31
traffic in women, 48·29
Web page defacement, 15·28, 15·30
Internet Assigned Numbers Authority (IANA), 30·34
Internet Control Message Protocol (ICMP), 5·16, 5·24–5·25, 18·8, 18·14, 18·20, 30·30
Internet Corporation for Assigned Names and Numbers (ICANN), 31·7
Internet Engineering Task Force (IETF), 7·28
PKI for X.509 Certificate (PKIX) working group, 37·8. *See also* X.509 certificate format
RFC 822, 37·18
RFC 1918, 16·13, 30·34
RFC 2196, 25·2
RFC 2196, *Site Security Handbook,* 44·8–44·9
RFC 2246, 7·29
RFC 2267, 16·13
RFC 2385, 5·23
RFC 2527, 37·17
RFC 2535, 5·26
RFC 2560, 37·21

Internet Engineering Task Force (IETF) (*Continued*)
RFC 2634, 5·28
RFC 2822, 5·28
RFC 3704, 16·13
RFC 3833, 5·26
RFC 3850, 5·28
RFC 3851, 5·28
RFC 4033–4035, 5·26
RFC 5280, 37·5, 37·7
Simple Network Management Protocol. *See* Simple Network Management Protocol (SNMP)
and TCP security, 5·23
Transport Layer Security. *See* Transport Layer Security (TLS)
Internet Exploder, 17·8
Internet Explorer (IE), 1·10, 7·27, 7·28–7·29, 18·24, 18·25
Internet Information Services (IIS), 18·8, 18·23, 18·24
Internet Key Exchange (IKE), 32·4–32·5
Internet Message Access Protocol (IMAP), 5·28
Internet Protocol (IP). *See also* TCP/IP (Transmission Control Protocol/Internet Protocol)
address, 5·3, 5·16, 30·34
and denial-of-service attacks, 18·2. *See also* Denial-of-service attacks (DoS)
function of, 5·14
IP Version 4 (IPv4). *See* IPv4
IP Version 6, 5·16–5·17
and layered standards architectures, 5·11
and protecting Web applications, 21·6
and security, 5·10
spoofing, 26·5, 31·10–31·11
Internet Protocol Telephony (IPT). *See* Voice over Internet Protocol (VoIP)
Internet Protocol Version 6. *See* IPv6
Internet Relay Chat (IRC) bots, 16·8–16·9
Internet Security Scanner (ISS), 46·3
Internet service providers (ISPs):
actions against denial-of-service attacks, 18·27–18·28
copyright issues, 11·22. *See also* Copyright law
and denial-of-service (DoS) attacks, 18·2
e-mail privacy issues, 11·32
and network access points, 5·8
parental tools for blocking Web content, 31·10
preventing abuses of anonymity and pseudonymity, 70·26
and security incidents, 56·11
and spam, 20·20–20·21
spam filtering, 20·22–20·23
Web hosting, 48·5
Internet Worm, 18·7
Interviews:
computer crime research methods, 10·9
Intimidation, 19·6
Intranets, 5·7
Intrusions:
alarms. *See* Alarms
analysis schemes, 27·5–27·6, 27·8–27·10
blended attacks, 55·13
defined, 27·2
detection, 26·11–26·12, 26·14, 26·18, 27·2–27·4, 41·9
e-commerce security services, 30·6
intrusion detection systems (IDSs), 27·2, 39·20
intrusion prevention systems (IPSs), 39·20, 41·9
monitoring, 27·5–27·8
overview, 27·4–27·6, 27·16
prevention, 26·11–26·13, 27·3, 27·6
process of and cyberinvestigation, 55·13–55·14
product selection and needs assessment, 27·13–27·16
reporting, 61·3–61·5
response, 26·12–26·13, 27·5, 27·10–27·13
wireless intrusion detection and prevention systems (WIDPS), 27·14
Invalid file access attempts log record, 53·13
Invalid logon attempts log record, 53·13
Inventory, and vulnerability management, 46·2
IP Traceback, 18·29–18·30
iPad, 17·2
iPhone, 17·2
iPod, 17·2
IPsec, 5·9, 5·17, 5·23, 25·4–25·5, 32·4, 32·11–32·12
IPSec Developers Forum, 25·5
IPv4, 5·14–5·16, 21·7, 26·16, 30·31, 30·34
IPv6, 21·7, 26·16–26·17, 30·31, 32·14–32·15, 32·19
ISACA, 74·19–74·20
Islamic countries, censorship in, 72·7–72·11
IT Baseline Protection Manual, 44·9
IT-Grundschutz Catalogues, 44·9
IT-Grundschutzhandbuch, 44·9
ITSEC. *See* Information Technology Security Evaluation Criteria (ITSEC)
iTunes Music Store, 42·8, 42·10

Jacket, 6·11
Java:
applets, 16·11, 17·3, 17·11–17·12, 17·13, 21·8, 48·35
buffer overflows, 39·13
and component-based software, 21·14
and content filtering, 41·10
and denial-of-service attacks, 18·13
and e-commerce vulnerabilities, 21·4, 21·22
and firewalls, 30·35
Java 2 Enterprise Edition (J2EE), 21·13
security features, 38·9
source language programs, 47·3

Java Run Time Environment, 17·11
Java Virtual Machine (JVM), 17·11
JavaScript, 16·11, 17·2, 17·3, 21·15, 41·10
JavaScript/ECMA Script, 17·15
JavaServer Pages, 15·30
Jerusalem virus (Friday the 13th virus), 2·11
Job Control Language (JCL), 47·4
Job scheduling, 53·17
Jobs, Steve, 1·8, 2·8
Joint application development (JAD), 39·7
Just-in-time (JIT) production, 30·24

Kant's categorical imperative, 70·25–70·26
KENBACK–1, 1·8
Kerberos, 3·12, 25·12, 28·10, 32·6, 36·9–36·10
Kerckhoff's Principle, 7·44
KERMIT, 30·34
Kernel mode, 24·9–24·10, 38·6–38·7, 41·2
Kernel panic, 18·9–18·10
Key encapsulation, 37·23
Key management system, 33·8–33·9
Key escrow, 37·23–37·24
Key recovery agent (KRA), 37·23
Key recovery practices statement (KRPS), 37·24
Key recovery requester (KRR), 37·24
Key Usage extension, 37·6
Keylogger Trojans, 2·13–2·14
Keyspace, 7·3, 7·10, 7·21, 15·16–15·17
Keystroke loggers, 15·15, 16·8–16·9
Kismet, 33·45
Kumho Tire Co. v. Carmichael, 73·3

L0phtCrack, 15·25, 25·9
L0pht Heavy Industries, 2·27
LANalyzer, 25·4
Land attack, 18·10
Language translation sites, 31·12–31·13
Laptops, 1·19, 36·12–36·13, 57·23
Lavasoft, 21·9, 48·15, 48·43
Law enforcement, cooperation with:
  and e-commerce, 30·20
  evidence handling, 61·7–61·12
  FBI. *See* Federal Bureau of Investigation (FBI)
  information sharing, 61·12–61·18
  memorandum of agreement, 61·8
  monitoring technologies, 69·14–69·15
  overview, 61·1–61·2, 61·17–61·18
  planning ahead, 61·3
  relevant laws, 61·3
  reporting, 61·3–61·7
  resources, 61·7–61·8
  security incidents, 56·11, 56·24–56·25
Layer 1 attacks, 33·23–33·24
Layered standards architectures, 5·11–5·12
LDAP, 32·6
Learning bridges, 6·25
Leased lines, 4·15, 5·5, 15·9–15·10
Least privilege, 24·4
Legal and regulatory compliance. *See also* Standards
  all-hazard mitigation plan, 23·43–23·44
  code security, 38·5, 39·10
  data classification, 67·3–67·6
  database encryption, 36·13–36·15
  disaster recovery, 59·16
  e-commerce, 30·43–30·44
  overview, 54·2, 64·1–64·2
  privacy, 60·18–60·23. *See also* Privacy law
  Voice over Internet Protocol (VoIP), 34·2–34·3
Legion of Doom, 2·24–2·25
Legion of Underground (LoU), 2·23
Levin, Vladimir, 2·10
Liability:
  and cyberinvestigations, 61·9
  destruction of e-mail records, 48·36
  employment termination, 45·15–45·16
  hostile work environment, 48·33–48·36
  identity theft, 60·18
  illegal copies of software, music, and videos, 48·31
  libel, 48·31
  negligent hiring and retention of employees, 45·2
  and physical security, 22·6–22·7, 22·11
  plagiarism, 48·32–48·33
  Website management, 30·43–30·44
Libel, 48·31
Licenses and licensing:
  Berkeley Software Distribution (BSD) License, 11·34
  and first sale doctrine, 11·9–11·10
  General Public License (GPL), 11·33–11·34
  Massachusetts Institute of Technology (MIT) License, 11·34
  open source code, 11·33–11·34
  shrink-wrap and click-wrap licenses, 11·4, 11·17
  and TRIPS anticompetitive restrictions, 11·38
Light-emitting diode (LED), 6·10
Lightweight Access Point Protocol (LWAPP), 33·41
Lightweight Directory Access Protocol (LDAP), 26·14, 33·7, 34·8, 37·17
Limewire, 16·8
LinkedIn, 48·5
Linux, 6·28, 18·2, 18·10, 18·21, 18·22, 25·9, 25·12, 25·14
Linux-wlan, 33·45
Liquids as security hazard, 3·17
List linking, 18·7
List servers, 48·23
Living organisms as security hazard, 3·18
LLC standard, 6·23–6·24
Load balancing, 26·24
Load condition errors, 38·12, 39·9

Local area networks (LANs):
access ports, disabling, 16·14
background, 1·10–1·11
characteristics, 6·2
components of, 6·2–6·3
Ethernet, 6·25
infrastructure security, 25·3–25·9
interconnection devices, 6·24–6·27
media, 6·2–6·3, 6·8–6·13
media access control (MAC) standard, 6·3, 6·13–6·15
and network-attached storage (NAS), 57·4
network control, 6·4
network interface card (NIC), 6·3
network operating systems, 6·3, 6·27–6·28, 25·9–25·16
overview, 5·4–5·5, 6·2
packet sniffing, 15·10–15·11
physical site security, 25·3
policy and procedure issues, 25·1–25·3
promiscuous mode, 25·4
security, generally, 25·1, 25·16
sniffers, 25·4. *See also* Packet sniffers
technology parameters, overview, 6·3
wireless (WLAN), 6·12, 15·12, 25·7–25·9
Lockdowns, 20·35
Locks and door hardware, 23·12–23·13, 36·3
Logic flow errors, 38·11, 39·9
Logical domain addresses, 21·7
Logical Link Control (LLC), 6·16
Logical Link Control Working Group, 6·17
Logical separation of voice and data, 34·7
Login:
information, trapping, 15·14
speed, 15·17
visitors, 47·6–47·7
Logs and logging:
centralized data stores, 53·21
classes, 53·17–53·20
exception reports, 53·22
file systems, 53·20
filtered queries, 53·21
and gateway security devices, 26·23, 26·28–26·29
log files, 30·44, 47·17–47·18, 53·22
log management, 53·11–53·17
log records, 61·11–61·12
log review, 46·4
log server, 39·20
and new versions of software, 47·7–47·9
record security, 53·15–53·17
records, analyzing, 53·21–53·22
transaction logs, 53·11
Longitudinal redundancy checks, 4·4–4·5
Logic bombs, 2·11, 16·4–16·5, 45·9
Lotus Notes, 7·27
"Love Bug" worm, 2·20, 41·11
Love Letter (ILOVEYOU) worm, 18·4, 25·11
Low-tech social engineering attacks, 19·4, 19·6–19·8

MAC protocol data unit (MPDU), 33·18
MAC service data unit (MSDU), 33·18
MacOS, 25·9, 25·15–25·16
Machine check interrupts, 4·7
Macro viruses, 16·4
MafiaBoy, 2·22, 18·15
Magic Quadrant analyses, 51·21–51·22
Magnetic discs, 36·1
Maintenance and repair:
personal computers, 4·24–4·25
and protection of infrastructure, 23·9–23·10
Website maintenance and updating, 30·32
Malicious code, 14·22, 48·13–48·14. *See also* Malware; Mobile code
Malicious threats, 13·5–13·6
Malware. *See also* Mobile code
in the 1990s, 2·17–2·18
antimalware, 26·16, 66·10
and antivirus technology. *See* Antivirus programs
automated malware attacks, 55·13
detecting, 16·11–16·12
and e-commerce, 21·9–21·10
early outbtreaks, 2·15–2·16
financial gain, 41·2
IRC bots, 16·8–16·9
overview, 16·1–16·2
prevention, 16·12–16·14
rootkits, 16·10. *See also* Rootkits
and social engineering attacks, 19·9–19·10
spyware, 16·9–16·10. *See also* Spyware
and system penetration, 15·31–15·32
threat model, 16·2–16·3
Trojans, 16·8–16·9. *See also* Trojan horses
viruses, 16·4–16·7. *See also* Viruses
worms, 16·7–16·8. *See also* Worms
Man-in-the-middle attacks, 21·12, 34·6, 36·6–36·7, 36·10, 37·4, 47·16
Managed Security Service Provider (MSSP), 26·33–26·34
Management:
audits, 63·22
controlling, 63·25
due diligence of care, 63·20
employee management, 45·3–45·12
failures, 63·16–63·18
functions, 63·23–63·25
integrating, 63·24–63·25
judgment and adaptation, 63·15
liability issues, 63·19–63·22
organizing, 63·23–63·24
planning, 63·23
policy, 63·12
responsibilities, 63·10–63·18
risk management. *See* Risk management

role of, 63·1–63·2
security administration, 63·25–63·28
sensitization, 66·11–66·12
SOX compliance, perspective on, 64·5–64·6
staffing, 63·25–63·26
stakeholders, 63·19–63·20
Veterans Affairs case study, 63·5–63·10
Management frame spoofing, 33·26–33·27
Management Information Bases (MIBs), 25·10
Management interfaces, 36·6–36·7
Mandatory access controls, 9·2, 9·6, 9·9
Mantraps, 23·24
Masonry partitions, 23·25
Massachusetts Data Security Regulations, 60·22
Massachusetts Institute of Technology (MIT) License, 11·34
Maximum segment size (MSS), 5·21
Media access control (MAC) standard, 6·3, 6·13–6·15, 25·4
Media Independent Handover Services Working Group, 6·19
Medical devices security initiatives, 71·50–71·51
Medical emergencies, 22·8, 22·23–22·24, 23·43
Medical records. *See* Healthcare security and privacy
Melissa virus, 2·18–2·19, 15·37, 18·4, 25·11
Memorandum of agreement, 61·8
Memory:
consumption, log records for, 53·15
main memory, 4·8
protection of, 24·5–24·11
read-only memory, 4·8–4·9
Memory dumps, 53·9–53·11
Message Application Program Interface (MAPI), 18·24
Message Authentication Code (MAC), 7·4, 7·30, 38·8
Meta-analysis, 10·11
Metacharacters, 15·30
Metadata, 4·4
Metatags, 42·14
Metropolitan Area Network (MAN) Working Group, 6·17
METT-TC analysis, 47·2
Michelangelo virus, 2·11
Microprocessors, 1·8
Microsoft Corporation:
antitrust trial, 19·6–19·7
Assistance Markup Language, 21·18
Authenticode, 17·6
Common Object Model (COM), 21·13
Data Classification Toolkit, 67·8
Distributed COM (DCOM), 21·13
Internet Explorer (IE), 1·10, 7·27, 7·28–7·29, 16·6
Internet Information Services (IIS), 18·8, 18·23, 18·24
Microsoft SQL Server 2005, 36·13–36·14
Outlook Express, 20·4–20·5
PCI DSS Compliance Planning Guide, 51·6
Point-to-Point Encryption, 25·8
Secure Development Lifecycle, 51·6, 51·13
software registration and antipiracy programs, 42·4–42·5
Trustworthy Computing initiative, 25·12
Windows. *See* Microsoft Windows
Windows Volume Shadow Copy Service (VSS), 57·19
Microsoft networking, 6·28
Microsoft Windows:
and ActiveX controls, 17·13–17·14. *See also* ActiveX
Challenge Handshake Authentication Protocol (MS-CHAP), 33·11–33·12, 33·33–33·34
and Help documents, 21·18
Microsoft Windows Server, 6·28
and network operating systems, 25·9
Windows 9x, 25·10
Windows 95, 25·10
Windows 98, 25·10
Windows 2000, 18·12, 24·14–24·20
Windows Defender, 25·12
Windows ME, 25·10
Windows NT Resource Kit, 25·13
Windows NT servers, 18·4
Windows Service Hardening, 25·12
Windows Vista, 16·13, 21·7, 21·18, 25·12, 42·5
Windows Update ActiveX control, 17·14
Windows XP, 25·11–25·12
Microsoft Word, 48·20
Microwave LANs, 6·12
mile2, 74·22
Miliefsky, Gary S., 1·17
Military:
information categories, 24·2, 24·12
information warfare, 14·5–14·8
integration of cyberwarfare into military planning, 14·27–14·28
and intrusion detection systems, history of, 27·4
operations security and acronyms, 47·2
spread spectrum radio transmission, 5·13
MIME. *See* Multipurpose Internet Mail Extensions (MIME)
MIS Training Institute (MISTI), 56·36
Misrepresentation, 3·2, 3·15, 15·3–15·6
Misuse of or failure to use information, 3·2, 3·16
Mitigation:
adaptive threat mitigation, 26·12
all-hazard mitigation plan, 23·43–23·44
collaboration tools, security breach prevention and mitigation, 35·20–35·22
IT systems interruption, 62·21
peer-to-peer (P2P) networking, security breach prevention and mitigation, 35·7–35·8
of risk, 62·10–62·16, 62·24

Mitigation (*Continued*)
short message service (SMS), security breach prevention and mitigation, 35·15–35·18
State and Local Mitigation Planning guides, 23·43
violence, prevention and mitigation, 23·7–23·8
Mitnick, Kevin, 2·5–2·6, 3·6, 15·6–15·7, 19·3
MITS Altair 8800, 1·8
Mobile Broadband Wireless Access (MBWA) Working Group, 6·19
Mobile code:
ActiveX, 17·7–17·13, 21·8, 30·35. *See also* ActiveX
client responsibilities, 17·14–17·15
defined, 17·3
and firewalls, 30·35
Java, 17·3, 17·11–17·13
malicious, 16·11, 17·3
misappropriation and subversion, 17·13
motivation and goals of malware, 17·5
as multidimensional threat, 17·13–17·14
overview, 17·1–17·3
restricted operating environments, 17·11–17·12
server responsibilities, 17·15–17·16
signed code, 17·6–17·11, 17·12–17·14
trust issues, 17·12–17·13
and Web servers, 17·2, 17·15–17·16
from World Wide Web, 17·3–17·5
Mobile devices. *See also* Cellular phones and modems; Wireless networks
backups, 57·23–57·24
intrusion detection, 41·13
and location privacy, 69·26–69·27
and software patches, 40·5
and virtual private networks, 32·9
Mobile computing, 53·27
Mobile IPv6, 26·17
Modal dispersion, 6·11
Models of computer security:
access-control matrix model, 9·3–9·5
Bell-LaPadula model, 9·2, 9·9–9·12, 9·18–9·19
Biba's strict integrity policy model, 9·2, 9·9, 9·12–9·14, 9·18–9·19
Chinese Wall model (Brewer-Nash model), 9·2, 9·17–9·18
Clark-Wilson model, 9·2, 9·9, 9·14–9·16, 9·18–9·19
Clinical Information Systems Security model, 9·18
and controls, 9·6–9·9
deducibility security, 9·19
discretionary access controls, 9·2, 9·6, 9·9
importance of models, 9·1–9·3
mandatory access controls, 9·2, 9·6, 9·9
noninterference security, 9·19
originator-controlled access control, 9·2, 9·6–9·7, 9·9
overview, 9·2
role-based access controls, 9·2, 9·7–9·9
terminology, 9·3
traducement, 9·18
typed access control model, 9·6
Modems, 1·10, 4·14–4·15, 5·4, 25·5–25·7
Modification of data, 3·2, 3·15
Monitoring. *See also* Monitoring and control (M&C) systems
access controls, 53·20
customers, 30·42–30·43
e-commerce security services, 30·6
environmental concerns, 53·18
file systems, 53·20
intrusion detection, 27·5–27·8
output quality, 47·13
passive network monitoring, 40·8
performance, 47·11–47·12
process activities, 53·19–53·20
resources, 47·12–47·13
system components, 53·19
system state, 53·18
Monitoring and control (M&C) systems:
automation, 53·7–53·8, 53·15
challenges, 53·25–53·29
change management, 53·24–53·25
checklist, 53·29–53·30
configuration protection, 53·25
control loop, 53·5
controlling vs. monitoring, 53·3–53·5
data aggregation and reduction, 53·21–53·22
distributed model, 53·7
human–machine interface (HMI), 53·7–53·8
internal systems, 53·6
job level, 53·15
log management, 53·11–53·17. *See also* Logs and logging
logging classes, 53·17–53·20
memory dumps, 53·9–53·11
notifications and reporting, 53·22–53·24
one-to-one model, 53·6–53·7
overview, 53·2–53·4, 53·29
performance considerations, 53·25
prevention, detection, and response as purpose of, 53·3
real-time monitoring, 53·8–53·9
resource allocation, 53·15
scope, defining, 53·5–53·6
snapshots, 53·8–53·9
system level, 53·15
system models, 53·6–53·11
system requirements, defining, 53·5–53·6
Monoalphabetical substitution ciphers, 7·11–7·12
Monte Carlo simulation, 62·23
Morris, Robert T., 2·16–2·17, 3·13, 16·7, 18·3, 18·5. *See also* Morris Worm
Morris Worm, 2·16–2·17, 16·7, 18·3, 18·5, 30·39, 56·4–56·5

Motion Picture Association of America (MPAA), 42·7
Motivation, 4·21
Movement as security hazard, 3·18
MP3 music files, 42·7–42·8
MPEG compression, 42·9, 42·11
MS-CHAP-v2, 33·33–33·34
Multicast listener discovery (MLD), 26·17
Multilevel security, 24·12
Multimode fiber (MMF), 6·11
Multiprotocol layer switching (MPLS), 32·7–32·8
Multipurpose Internet Mail Extensions (MIME), 5·28, 37·5
Multistation access units (MAUs), 6·25
Music downloads, 42·6–42·8, 42·10, 48·31
MyMP3.com, 42·7
MySpace, 16·6, 48·5, 69·27

N-Tier architecture, 32·16
Napster, 11·23–11·24, 35·4, 42·7–42·8
NAT-Traversal, 34·8–34·9
National Bureau of Standards (NBS), 7·20
National Commission for Certifyhing Agencies (NCCA), 74·9
National Computer Security Center (NCSC):
  Orange Book, 1·14–1·15, 17·11, 17·15, 25·11, 51·10
  system ratings, 25·11
National Electric Code (NEC), 23·8
National Institute of Standards and Technology (NIST):
  approved DES implementations, 7·28
  Computer Security Incident Handling Guide, 56·34
  and development of the common language, 8·3
  Federal Information Processing Standards (FIPS) 51·6, 51·23, 54·9
  FIPS Publication 199, 54·9
  FIPS Publication 200, 54·10
  and quantum cryptography, 7·43
  risk framework, 54·10
  SP 800-16, "Information Security Training Requirements," 49·17–49·18
  SP 800-16 REV1, 49·3
  SP 800-26, 54·9
  SP 800-30, 54·10
  SP 800-37, 54·10
  SP 800-48, "Wireless Network Security," 35·17
  SP 800-53, 40·2, 54·8–54·9, 54·10, 71·45–71·46
  SP 800-55, 40·10
  SP 800-60, 54·4, 54·10
  SP 800-78-3, 37·22
  SP 800 series, 38·7
National Incident Management System (NIMS), 23·2–23·3
National Infrastructure Protection Plan (NIPP), 23·3
National Organization for Competency Assurance (NOCA), 74·9
National Response Framework (NRF), 23·3
National Security Agency (NSA), 7·3, 7·2069·2–69·3, 70·20
  communications-monitoring activities, 31·4
  Computer Security Center, 24·14
  and elliptic curve cryptography, 7·36
  and Microsoft Vista, 25·12–25·13
  Rainbow Series, 1·14–1·15
  Secure Hash Algorithms (SHA), 34·9
  Security Guidelines Handbook, 44·6–44·7
  TEMPEST compliance, 25·5
National Vulnerability Database, 18·8
Netscape Navigator, 7·27, 7·28–7·29
Netstumbler, 33·45
Network access control (NAC), 32·9
Network access points (NAPs), 5·8
Network Access Protection, 25·12
Network activity log records, 53·14
Network address translation (NAT), 5·3, 26·11, 26·17, 32·5, 34·8–34·9
Network anomaly detection (NAD), 16·14
Network Associates, 25·4
Network attached storage (NAS), 36·4, 36·5
Network connectivity, 53·17–53·18
Network File System (NFS), 36·8–36·9, 57·4
Network interface card (NIC), 5·3, 6·3, 25·4
Network intrusion prevention systems (NIPS), 26·18, 53·17
Network Monitor, 25·4
Network monitoring, 16·14, 34·8
Network operating systems (NOS), 1·10, 6·3, 25·9–25·16
Network proxy, 42·14
Network security:
  adaptive threat mitigation, 26·12
  added features, 26·26–26·27
  allowed paths, 26·9–26·10, 26·21
  application inspection, 26·12–26·13
  auditability, 26·29
  complex protocols, 26·14–26·15
  connection termination, 26·11–26·12
  content control 26·15
  cost, 26·31–26·32
  data leakage prevention, 26·15
  device evaluation, 26·25–26·34
  encryption, 26·13–26·14
  features, 26·30
  general capabilities, 26·4
  GSD positioning, 26·19–26·20
  host protection, 26·17
  identity-based enforcement, 26·14
  infrastructure requirements, 26·25
  learning, 26·30
  maintenance, 26·22–26·24
  Managed Security Service Provider (MSSP), 26·33–26·34
  management, 26·27–26·28

Network security (*Continued*)
monitoring, 26·20–26·22
overview, 26·34
performance, 26·26
perimeter security, 26·16–26·17
policy, 26·21–26·22
proxy servers. *See* Proxy servers
system-level actions, 26·12
throughput, 26·26
tunneling, 26·10, 26·17
unified threat management (UTM), 26·4, 26·6
vendors, 26·32–26·3
Networks:
application standards, 5·28–5·29
monitoring for intrusion detection, 27·5–27·7
orthogonal, 16·12, 16·14
passive network monitoring, 40·8
protocol risks, 5·9, 21·10–21·12
scanners, 46·4–46·5
security, 66·9–66·10
simple home PC network, 5·2–5·4
standards, 5·9–5·14, 5·24–5·28
terminology, 5·2
Transmission Control Protocol (TCP). *See* TCP/IP (Transmission Control Protocol/Internet Protocol); TCP (Transmission Control Protocol); Transmission Control Protocol (TCP)
User Datagram Protocol (UDP), 5·24
NetZip, 48·14
Next-generation firewall (NGFW), 26·4, 26·6
Nigerian 419 fraud (advance-fee fraud), 2·21, 16·13
Noncompetition agreements, 45·16
Noninterference security, 9·19
Nonmalicious threats, 13·6–13·6
Nonrepudiation, 3·12–3·13, 7·5–7·6, 28·5, 37·6–37·7
Nonstandard IT components, 40·5
Norm of reciprocity, 50·19
North American Free Trade Agreement (NAFTA), 11·35
Novell, 25·4, 42·5
NetWare, 1·10, 6·28, 25·9
NRC System Security Study Committee, 1·14–1·17

Object Management Group (OMG), Common Object Request Broker Architecture (CORBA), 21·13
Observation, as means of information loss, 3·2, 3·16
Octet, 5·2, 5·20–5·21
OmniPeek, 33·46
Onion routing, 31·2, 31·12, 35·8
On-Line Certificate Status Protocol (OCSP), 37·20–37·21
Online auctions, 48·15, 48·41
Online dating, 48·28–48·29, 48·39–48·40
Online gambling, 48·15 48·26–48·27, 48·41–48·42
Online shopping, 48·14, 48·24–48·26, 48·40–48·41
Open design, 24·4
Open Shortest Path First (OSPF), 5·27
Open source code, 11·33–11·34
Open Systems Interconnection (OSI) Reference Model, 5·11, 6·15–6·17, 20·2
OpenVMS, 17·11
Operating system:
and e-commerce security, 21·21
file sharing, 24·11–24·12
fingerprinting, 15·23
formatting, and disposal of data, 57·31–57·32
interapplication conflicts, 38·12, 39·10
known-good boot medium, 47·13
memory protection, 24·4–24·6
mode of processor execution, 24·9–24·10
monitoring for intrusion detection, 27·7
multiuser, 17·11–17·12
new versions, 47·13
patches, 21·22, 47·13–47·14
protection mechanisms, 24·4–24·11
protection policies, types of, 24·1–24·2
restricted, 17·11–17·12
security kernel, 38·6–38·7
security requirements, 24·2–24·4
sharing resources, 24·4–24·5
Trojan horse defense, 24·14. *See also* Trojan horses
trusted systems, 24·12–24·14
virtualization, and software patches, 40·5
and Web applications, 21·7
and writing secure code, 38·6–38·7
Operation Copycat, 42·7
Operation Sundevil, 2·26
Operations security:
access to operations center, 47·5–47·6
cloud computing and production systems, 47·18–47·20
data protection, 47·14–47·16
data validation, 47·16–47·18
operating system, 47·13–47·14
operations defined, 47·2–47·3
operations management, 47·4–47·13
overview, 47·1–47·3, 47·20–47·21
Opportunistic key caching (OKC), 33·20
Opportunity, 4·21
Optical fiber, 5·13, 6·2, 6·10–6·11, 15·11, 25·5
Oracle, 19·6, 36·14
Orange Book, 1·14–1·15, 17·11, 17·15, 25·11, 51·10
Organizational culture, 15·2–15·3
Organization for Economic Cooperation and Development (OECD), 69·2, 69·8
Originator-controlled access control, 9·2, 9·6–9·7, 9·9

Orthogonal frequency-division multiplexing (OFDM), 6·12
Orthogonal networks, 16·12, 16·14
OS/360, 17·11
OSI layers, 5·11
Out-of-office hosts, 40·5
Output format errors, 38·15, 39·12
Output quality, monitoring, 47·13
Outsourcing:
  and application service providers (ASPs), 30·44–30·45
  and collaboration tools, 35·21–35·22
  consequences of, 68·10–68·11
  cost of, 68·13–68·14
  defined, 68·3
  distinction, 68·4
  and e-commerce, 30·26
  computer security incident response, 56·6–56·7
  environmental factors, 68·16
  and extranets, 32·17
  failure, reasons for, 68·8–68·19
  insourcing, 68·4–68·5
  intellectual property risks, 68·17–68·23
  international economics, 68·14–68·15
  intrusion detection and prevention, 27·15–27·16
  labor, 68·17
  nearshoring, 68·5
  offshoring, 68·5
  overview, 68·2–68·3, 68·30–31
  political issues, 68·15–68·16
  purpose and intent, articulating, 68·12–68·13
  reasons for, 68·6–68·8
  risk, nature of, 68·11–68·12
  security functions, 68·23–68·30
  social culture, 68·14
  travel, 68·16–68·17
Overflow, 4·6
Ozemail, 18·4

Pacific Telephone and Telegraph (PT&T), 2·5
Packet filtering, 26·5
Packet sniffers, 6·2, 15·10–15·11, 15·26, 16·11, 25·4, 69·14
Packet-switching networks, 15·10
Packets, 5·7–5·8
PairGain Technologies, 18·10
Pairwise master key (PMK), 33·9, 33·16, 33·20
Pairwise transient key (PTK), 33·14, 33·16
Pandemics, 22·14, 22·16. *See also* Physical threats
PANIX, 18·3, 18·5, 18·12
PAPA MS-Excel virus, 18·4
Parameter passing errors, 38·11–38·12, 39·9
Paravirtualization, 53·28
Parental tools for Web content filtering, 31·10, 31·13, 48·36
Paris Convention for the Protection of Industrial Properties, 11·19, 11·35
Parity, 4·4–4·5
Passfaces software, 28·13
Passive network monitoring, 40·8
Passwords:
  access to by system administrators, 28·5–28·6
  and authentication principles, 28·2–28·3, 28·5
  bypass passwords, 23·16
  changing, need for, 28·12
  cracking, 15·25–15·26, 25·9, 28·5–28·6, 28·9, 46·4
  and database security, 21·19–21·20
  dictionary attacks, 28·9–28·10, 28·16
  encryption, 28·9–28·11. *See also* Encryption
  failed attempts, 28·8–28·9
  guessing, 15·18, 28·8–28·9
  hashed, 28·9–28·10
  and local area networks, 6·2, 25·9
  MacOS, 25·15
  and nonrepudiation, 28·5
  one-time, 28·7, 28·17
  overview, 28·19–28·20
  Passfaces software, 28·13
  and Public Key Infrastructure (PKI), 28·8. *See also* Public Key Infrastructure (PKI)
  same password at multiple sites, 28·8
  and server spoofing, 28·11–28·12
  sharing, 28·7
  and smart cards, 28·15–28·16
  sniffing, 28·10–28·11
  and system penetration techniques, 15·16
  theft, 28·6–28·7
  and Trojan horses, 28·6
  zero-knowledge password proofs, 28·11
Patches:
  collaboration tools, 35·22
  management, 40·1–40·9
  operating systems, 21·22, 26·22, 47·13–47·14
Patent Cooperation Treaty (PCT), 11·35
Patent law:
  disclosure requirement, 11·19
  infringement, 11·19–11·20
  international, 11·19
  overview, 11·19, 42·2
  Patent Cooperation Treaty (PCT), 11·35
  patent trolls, 11·43–11·44
  and TRIPS, 11·37–11·38
Patent trolls, 11·43–11·44
Payment Card Industry Data Security Standards (PCI DSS), 21·8, 26·3, 30·12, 40·2, 53·25, 54·2, 54·7–54·8, 71·15
PC Cyborg Trojan, 2·13
PDF (portable document format), 44·14, 48·20
PEAP (protected EAP), 33·11–33·12
Pedophiles, 48·13
Peer-to-peer (P2P) networking:
  abusing company resources, 35·5
  and application security, 5·29

Peer-to-peer (P2P) networking (*Continued*)
BitTorrent, 35·8
business threats, 35·5–35·7
case study, 35·9
confidentiality, loss of, 35·6
illegal content, 35·4
and IP addresses, 5·26–5·27
Linux software distribution, 35·2, 35·8
and malware, 16·8
and music downloads, 42·6–42·8
Napster, 35·4, 42·7–42·8
overview, 35·2, 35·22–35·23
security breach prevention and mitigation, 35·7–35·8
security incident response, 35·8
uses of, 35·4–35·5
and video piracy, 42·8
Penetration testing:
collaboration tools, 35·22
firewalls and gateway security devices, 26·22
secure code, 38·15–38·16, 38·19–38·20
and security-awareness programs, 49·33–49·34
and vulnerability assessment, 46·4, 46·7–46·12
Performance:
errors, 38·14–38·15, 39·12
monitoring, 47·11–47·12
(speed) errors, 38·14–38·15
Perimeter zones, 26·18
Perl, 21·15, 21·20
Personal computers (PCs):
history of, 1·8–1·10
laptops, 1·19, 36·12–36·13
maintenance and repair, 4·24–4·25
and networks, 5·3
and productivity, 1·9
security issues, 4·20–4·25
spyware. *See* Spyware
Personal Digital Assistants (PDAs), 21·8
Personal identification number (PIN), 17·8, 28·3, 28·9, 28·15
Personality, theories of, 50·4–50·7
Personally identifiable information (PII), 1·19, 36·10
PestPatrol, 48·43
Pharming, 18·11, 19·9
Phishing, 2·21, 5·28, 15·32, 16·11, 17·4, 18·10, 19·9, 20·1–20·3, 20·27–20·31, 21·19
Phone-hacking scandal (U.K.), 69·3–69·5
Phone phreaking, 17·4
Photodiodes, 6·10
Phrack, 2·25, 15·39, 18·4
Physical access:
control, 1·9, 1·11, 16·3
off-hour visitors, 47·6–47·7
and social engineering, 19·5
visitor badges and log in, 47·6–47·7
Physical layer standards, 5·12–5·13
Physical losses, 3·10, 3·17–3·18, 4·11–4·13, 4·22. *See also* Power failures and disturbances
Physical site security:
alarms. *See* Alarms
access points, 23·23–23·24
and confidential design details, 23·6–23·7
electrical distribution panels, 23·24
electrical power issues, 23·28–23·30
emergency intercoms, 23·24
emergency power, 23·30–23·36
environmental control, 23·36–23·41
local area network, 25·3
mantraps, 23·24
premises security inspection, 23·22
protective construction, 23·24–23·27
remote spying devices, 23·42
and social engineering, 19·20
sprinkler systems, 23·22
surveillance systems, 23·8, 23·11, 23·18–21
violence, 23·42–23·43
wiretaps and bugs. *See* Wiretapping
Physical threats:
airborne particles, 22·21
assessment process, 22·9–22·14
bombs, 23·42–23·43
civil, political, and economic disruptions, 22·25
cleaning and maintenance, 22·23
confidential information regarding, 22·25–22·26
coordinated attacks, 22·25
costs, 22·5, 22·13–22·14
determining vulnerability to, 22·13
DHS/FEMA methodology, 22·7, 22·8–22·9
dirt and dust, 4·12
and e-commerce, 30·25
electromagnetic pulse attack (EMP), 22·20
emergency response plans, 22·8
fire and smoke, 22·20, 23·38–23·41
food and drink, 4·22, 23·10
hazardous material incidents, 22·20
health threats, 22·3, 22·16
high-energy radio-frequency (HERF) weapons, 22·19
humidity, 4·11–4·12, 22·22, 23·37
identifying, 22·10–22·11
illicit workstations, 22·24
leaks, 22·22
liability issues, 22·6–22·7
medical emergencies, 22·8, 22·23–22·24, 23·43
natural hazards, 22·15–22·16
off-hour visitors, 22·22–22·23
overview, 22·2–22·9
premises security, 22·5–22·6
and productivity, 22·4
radiation, 4·12
reporting, 61·5–61·6

rodents and insects, 23·11
solar activity, 22·25
sources of information and assistance, 22·11–22·12
storage room, 22·23
targets, 22·4
temperature, 22·22
temperature as security hazard, 3·17, 4·11
terminology, 22·7–22·8
terrorism. *See* Terrorism and terrorists
toxic threats, 22·20
uniform, comprehensive planning process, 22·8–22·9
utility disruptions, 22·25
violence, 23·42–23·43
water, 4·12
wiretaps. *See* Wiretapping
workplace violence, 22·3, 22·4–22·5, 22·21–22·22
Piggybacking, 19·8, 23·24, 61·12
PING, 30·27
Ping, 5·24
Ping of Death, 18·8, 25·15
Pinging, 5·24–5·25
PIPA. *See* Preventing Real Online Threats to Economic Creativity and Theft of Intellectual Property (PROTECT IP) Act (PIPA)
Piracy. *See also* Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS)
antipiracy techniques, hardware-based, 42·5–42·12
antipiracy techniques, software-based, 42·3–42·5
database protection, 11·21
and digital rights, 42·2–42·12
marketplace issues, 11·20–11·21
music downloads, 42·6–42·8, 42·10
overview, 11·20
policies on, 48·31
Stop Online Piracy Act (SOPA), 11·41–11·43
terminology, 42·19–42·22
types of, 42·2–42·3
PKI. *See* Public Key Infrastructure (PKI)
Plagiarism, 48·32–48·33
Plan-Do-Check-Act (PDCA) cycle, 54·13
Point-to-Point Tunneling Protocol (PPTP), 25·8
Pointer (PTR) record spoofing, 18·10
Policy and procedure:
collaboration tools, 35·20
e-commerce security services, 30·6
gateway security device policy, 26·21
information security, 66·12–66·13
local area networks, 25·1–25·3
peer-to-peer networking, 35·7
security incident response, 56·7–56·10
security policy, 44·1–44·15, 66·2–66·15
short message service (SMS), 35·16
writing secure code, 38·1–38·5
Polyalphabetical substitution ciphers, 7·12–7·13
Polymorphic viruses, 16·6
Ponzi schemes, 48·10–48·11
Pop-ups, 17·3
Pornography, 1·20, 48·13, 48·34–48·36, 48·38
Port Login (PLOGI), 36·8
Port numbers, 5·21–5·22
Port scanning, 15·22, 16·11, 46·4
Portable document format (PDF), 44·14, 48·20
Ports, disabling, 30·29
Positive-intrinsic-negative (PIN) photodiode, 6·10
Post Office Protocol (POP), 5·28
Poulsen, Kevin, 2·9
Power failures and disturbances, 4·11, 4·19–20, 4·22–4·23, 23·28–23·36, 59·2. *See also* Disaster recovery
Pre-shared key (PSK), 5·3, 33·14, 33·31–33·32
Preliminary evaluation phase, security policy development, 66·3–66·11
Pretty Good Privacy (PGP), 7·26, 37·2, 37·13
Preventing Real Online Threats to Economic Creativity and Theft of Intellectual Property (PROTECT IP) Act (PIPA), 11·40–11·41
Privacy. *See also* Confidentiality
and biometric authentication, 29·17, 29·19, 29·20–29·22
and digital rights management (DRM), 42·3
and e-commerce, 21·8
and electronic communications, 11·29–11·33
Electronic Communications Privacy Act (ECPA), 11·30
European approaches to, 69·8–69·12
insurance coverage for breaches of, 60·17–60·23
invasion of, remedies for, 11·29–11·30
laws, regulations, and agreements, 69·5–69·7
and online shopping, 48·24–48·25
and public directories, 21·18
Stored Wire and Electronic Communications and Transactional Records Act (SCA), 11·30, 11·32–11·33
and surveillance systems, 23·11, 23·20
terminology, 42·19–42·22
U.S. approaches to, 69·12–69·28
and Web monitoring, 31·12
Website privacy policies, 30·20
and Website security, 30·42–30·43
Wiretap Act, 11·30–11·32
workplace, 69·23–69·24
Privacy Act of 1974, 67·3, 69·13, 71·14
Privacy-enhancing technologies (PET), 31·2, 31·12, 42·14–42·15
Private keys, 7·5, 7·8, 7·23–7·28, 7·32–7·33, 7·44, 37·3, 37·5, 37·23
Privilege management:
access control privileges, 23·12
minimum necessary privilege, 17·12

Privilege management (*Continued*)
operating system security, 24·4, 24·9–24·10, 24·12, 24·16–24·20
Public Key Infrastructure, 37·25, 37·26
Procedural languages, 47·4
Procedural statements, 47·4
Procedures, 44·3. *See also* Policy and procedure
Process activities, monitoring, 53·19–53·20
Process flow, monitoring, 53·17
Process initiation log records, 53·12
Process termination log records, 53·12
Product assessment:
common testing, 51·31–51·32
consortium-based approaches, 51·16–51·18
hacking approaches, 51·19
open source model, 51·18–51·19
standards. *See* Standards
third-party commercial assessments, 51·20–51·22
trade press, 51·19–51·20
vendor self-declarations, 51·15
Product ciphers, 7·19–7·20
Production system, 47·2–47·3, 47·9, 47·18–47·20
Program check interrupts, 4·7
Program status word (PSW), 24·6
Programmer libraries, 47·3
Programmers, production program access, 47·14
Projectiles as security hazard, 3·18
Prosocial behavior, 50·16–50·17
Protection profiles (PPs), 51·29–51·31
Protocol tunneling, 31·11
Proximity cards, 23·13–23·14
Proxy servers:
monitoring, 53·19
network proxy, 42·14–42·15
and network security, 26·13, 26·16
and Web monitoring, 31·2, 31·9, 31·13
Psychology:
computer criminals. *See* Computer criminals
dangerous information technology insiders. *See* Insiders, information technology
and social engineering, 19·11–19·14
social psychology. *See* Social psychology
PTR record spoofing, 18·10
Public directories and private documents, 21·18
Public Key Infrastructure (PKI):
architecture, selecting, 37·13
background, 37·2–37·4
Certificate Policy (CP), 37·8–37·9, 37·17
Certificate Revocation List (CRL), 37·6, 37·13, 37·16–37·23
Certification Authority (CA), 17·6, 21·10–21·11, 37·5–37·22. *See also* Digital certificates
certification practice statement (CPS), 37·8–37·9
costs, 37·27
cross-certification, 37·14
and cross-domain authentication, 28·18
and EAP, 33·11
enterprise PKI, 37·7–37·8
interoperability, 37·14–37·18
key expiration and rekeying, 37·22
key recovery, 37·23–37·25
and mobile code, 17·6
need for, 37·4–37·5
and passwords, 28·8
privilege management, 37·25
proofing (vetting), 37·10
recertification, 37·22
registration authority (RA), 37·7–37·9
and remote access VPNs, 32·6
revocation, 37·6, 37·13, 37·16–32·22
and signed code, 17·6, 17·7
and soft tokens, 28·16
trust, levels of, 37·10
trust models, 37·11–37·13
trusted archival services, 37·25–37·27
trusted paths, 37·10–37·11
trusted time stamp, 37·26–37·27
Public keys:
encryption, 5·28, 7·5, 7·21–7·27, 7·37–7·38
key recovery, 37·23–37·25
and onion routing, 31·2
public key cryptosystems (PKCs), 37·2–37·4
rekeying, 37·22–37·23
and server spoofing, 28·11–28·12
smart cards, 28·15
and soft tokens, 28·16–28·17
validity period, 37·22–37·23
X.509 standard. *See* X.509 certificate format
Public libraries, censorship in, 72·19–72·21
Public relations, 56·32
Public schools, censorship in, 72·21
Punched-card system, 1·4
Python, 21·15

Quality assurance, software, 39·2–39·3, 47·10–47·12. *See also* Software development
Quality control, 47·10
Quality of service (QOS), 34·7
Quantum computing, 7·41–7·42
Quantum cryptography, 7·39–7·44
Quarantines, 20·35
Quicken, 17·8

Race condition errors, 38·12, 39·9, 39·9, 39·14, 48·23
Radiation, 4·12
Radio-frequency identification (RFID), 23·12, 23·13, 53·27
Radio Regulatory Technical Advisory Group (RR-TAG), 6·19
Radio signals, 5·4, 5·12
RADIUS, 32·6, 33·4, 33·25, 33·44
Radon, 22·20

Random access memory (RAM), 15·17–15·18, 36·2
Random Access Method of Accounting and Control (RAMAC), 1·5
Random sampling, 10·5–10·6, 10·8
Rapid application development (RAD), 39·7
Rapid Spanning Tree Protocol (RSTP), 5·13
Raw sockets, 25·12
Read-after-write, 4·6, 4·9
Read-only access, 36·9
Read-only file security, 30·39–30·40
Real-time Transport Protocol (RTP), 34·4
Recovery procedures, 4·20, 36·11, 57·6–57·7
Readers, specialized, 42·6–42·10
Recording Industry Association of America (RIAA), 42·6–42·8
Redundancy, as security element, 23·8–23·9
Redundancy checks, 4·4–4·5
Redundant Array of Independent Disks (RAID), 23·9, 36·4, 57·4–57·5
Reference monitor concept, 24·12–24·13
Registered Jack (RJ), 5·3–5·4
Registration authority (RA), 37·7–37·9
Registry keys, 17·14
Regression testing, 39·16, 39·18
Rekeying, 37·22–37·23
Reliability:
  protocols, 5·20
  UDP, 5·24
Remote Access Dial-In User Service (RADIUS), 25·8, 26·14
Reordering of data, 3·15
Replacement of data, 3·2, 3·15
Replication, 4·6
Repudiation, 3·12–3·13, 3·15–3·16, 28·5. *See also* Nonrepudiation
Reputation-based scanning, 41·9–41·10
Reputation, damage to, 48·2–48·12
Requirements analysis, 38·5–38·6, 39·6
Research methodology, computer crime, 10·3–10·11
Resilient Packet Ring (RPR) Working Group, 6·18–6·19
Resource exhaustion, 38·12, 39·9–39·10
Resource overload, 40·6
Resource utilization log record, 53·14
Resource starvation, 18·12–18·13
Responsibility matrix, 45·10–45·11
Retinal or iris scanners, 23·15
Reverse engineering, 42·17
Reverse-path filtering, 31·11
Reverse social engineering, 19·11
Revocation:
  Authority Revocation List (ARL), 37·19–37·20
  Certificate Revocation List (CRL), 37·6, 37·13, 37·16–32·22
  public key, 37·18–37·22
Rich text format (RTF), 44·14, 48·20
Right to Financial Privacy Act of 1978, 69·14
Ring topology, 6·4–6·5
Risk, definition of, 1·2
Risk assessment:
  algorithms, 62·17–62·18
  e-commerce Websites, 30·22–30·25
  loss potential, 62·18–62·21
  objective, 62·4–62·5
  questionnaires, use of, 62·6
  return on investment (ROI), 62·5–62·6, 62·11, 62·13–62·16
  risk event parameters, 62·21–62·22
  risk model, 62·6–62·10
  sensitivity testing, 62·23–62·24
  techniques, 62·16–62·24
  Threat Effect Factors, 62·22–62·23
Risk management. *See also* Risk assessment
  annualized loss expectancy (ALE), 62·3–62·18
  overview, 62·2–62·3, 63·3
  legal and regulatory compliance, 62·3–62·4
  mitigation of risk, 62·10–62·16
  risk classification, 63·18–63·19
  risk defined, 62·2
  standards, 62·3–62·4
Rivest Cipher 4 (RC4), 25·8
RJ-11 connectors, 5·3
RJ-45 connectors, 5·3
Robert T. Stafford Disaster Relief and Emergency Assistance Act of 1988 (Public Law 100-707), 23·3
Robust Security Network Associations (RSNAs), 33·9, 33·34–33·35
Robust Security Network (RSN), 33·4, 33·8, 33·9–33·20
Rogue access points, 33·21–33·22, 33·39, 33·41
Role-based access controls, 9·2, 9·7–9·9
Root capability, 21·5
Rootkits, 8·14, 15·26, 16·10, 21·4, 41·2, 55·14
Routers:
  and network interconnection, 6·26–6·27
  access router, 5·3
  denial-of-service attacks (DoS), 18·13
  dynamic routing protocols, 5·27
  and layered defense, 16·13
  and network security, 16·13
  onion routers, 31·2, 31·12, 35·8
  terminology, 5·7
Routing Information Protocol (RIP), 5·27
RSA Data Security Company:
  CSIRT management conferences, 56·36
  password generators, 28·17
  RSA algorithm, 7·25–7·27, 7·36–7·38, 37·16, 37·22–37·23
  SecurID, 28·17, 28·18
RST (Reset + Restart) message, 5·9, 5·19–5·20
Russia, censorship in, 72·11–72·12
Russian Business Network (RBN), 2·27–2·28

S/HTTP, 21·10
S/MIME. *See* Secure/Multipurpose Internal Mail Extensions (S/MIME)
*-property (star property), 24·12, 24·14
Sabotage, 2·3–2·5
Salami fraud, 2·10–2·11
Salience effect, 19·13
Sandboxes, 17·5, 17·11
SANS Institute. *See* System Administration and Network Security (SANS) Institute
Sarbanes-Oxley Act of 2002 (SOX):
  audit and certification, 54·25
  and code security, 38·5, 39·10
  compliance, 54·21
  control framework, 54·21
  and data classification, 67·5
  and documentation, 39·16, 39·20
  internal control weaknesses, 38·5, 39·10
  management perspective on, 64·5–64·6
  and need for network-level centralized control, 26·3
  overview, 64·2–64·4
  provisions of (list of titles), 64·3
  regulatory compliance, 34·2–34·3
Section 404, 54·20–54·21, 64·4–64·5
  and security audits, 54·2, 54·4, 54·20–54·26
  and security planning, 23·3
  and VoIP compliance, 34·2
SATAN (Security Analysis Tool for Auditing Networks), 15·21–15·22, 46·3
SCADA. *See* Supervisory control and data acquisition (SCADA)
Scanners, 15·22–15·27, 41·4–41·6, 46·4–46·5
Schneider, Jerry Neal, 2·5
Schools. *See also* Colleges and universities
  Web monitoring and content filtering, 31·3, 31·12
Scrambler, 2·12
Script kiddies, 16·2, 18·2, 33·40, 55·13
Scripting languages, 21·15
Search engines, 19·8, 31·7, 31·13
Secure Digital Media Initiative (SDMI), 42·12
Secure/Multipurpose Internal Mail Extensions (S/MIME), 5·28, 20·19, 21·10, 37·7, 37·16
Secure Real-Time Protocol, 34·9
Secure remote access, 32·3
Secure Session Initiation Protocol (SSIP), 34·4, 34·9
Secure Shell (SSH), 5·28, 25·4–25·5, 28·10, 31·11
Secure Sockets Layer (SSL). *See also* Transport Layer Security (TLS)
  application standards, 5·28
  and e-commerce, 21·10, 30·11–30·12
  and encryption, 7·28–7·30
  and mobile code, 17·4
  and password encryption, 28·10–28·11
  SSL/TLS. *See* Secure Sockets Layer/Transport Layer Security (SSL/TLS)
  and trust chain, 37·5
Secure Sockets Layer/Transport Layer Security (SSL/TLS), 5·28, 30·41, 32·4, 32·7
SecurID, 28·17, 28·18
Securities and Exchange Commission (SEC), 60·20–60·21
Security, definition of, 1·2
Security administrators, 47·4–47·5
Security Assertion Markup Language (SAML), 28·15–28·16, 32·19
Security descriptors, 24·17–24·20
Security doors, 23·25
Security Guidelines Handbook, 44·6–44·7
Security ID (SID), 24·16–24·20
Security incident and event management (SIEM) system, 26·23
Security incidents:
  attacks, 8·11–8·16
  common language for, 8·1–8·20
  defined, 8·1
  events, 8·4–8·11
  incident defined, 8·15
  overview, 56·2–56·3
Security kernel database, 24·12–24·13
Security officer, 47·4–47·5
Security through obscurity, 5·6
Security planning:
  all-hazard mitigation plan, 23·43–23·44
  cost-benefit analysis, 23·44–23·45
  implementation, accountability, and follow-up, 23·46
  security response plan, 23·45–23·46
Security policy:
  antimalware measures, 66·10
  archives, 66·10–66·11
  backups, 66·10–66·11
  business resumption planning, 66·11
  collaboration in developing, 66·2–66·3
  computer operations, 66·8
  data access controls, 66·9
  data destruction, 66·10–66·11
  development phases, 66·3–66·14
  disaster recovery, 66·11
  explaining, 50·12–50·13
  implementation, 66·13–66·14
  incident response, 66·11
  maintenance, 44·14–44·15, 66·14
  management sensitization, 66·11–66·12
  need for, 66·2
  needs analysis, 66·12
  network and communications, 66·9–66·10
  organization, 44·11–44·12
  overview, 66·14
  physical security, 66·7
  policies and procedures, 66·12–66·13
  policy development group, 66·3
  preliminary evaluation phase of policy development, 66·3–66·11
  publishing, 44·12–44·14
  recommendations for creating and implementing, 44·15

resources for policy writers, 44·3
review of, 44·15
software development, 66·7–66·8
standards, 44·3–44·10
templates, 44·9–44·10
terminology, 44·2–44·3
updating, 44·15
and Web application systems, 21·5–21·6
writing, 44·10–44·11
Security response plan, 23·45–23·46
Security services, 30·5–30·13
Security target profiles (STs), 51·31
Security University, 74·22–74·23
Security Working Group, 6·18
Seduction, 19·5–19·6
Seeding, 39·15
SEEK, 7·27
Self-Monitoring, Analysis, and Reporting Technology (SMART), 4·10
Semiconductor Chip Protection Act of 1984 (SCPA), 11·12
Sendmail, 16·7, 17·16
Sensitive compartmented information facilities (SCIFs), 15·14
Separation of duties, 45·8–45·10, 47·4–47·5
September 11, 2001, attacks, 4·18, 22·2, 30·19, 59·3, 59·5, 62·12. *See also* Terrorism and terrorists
Serial broadcast, 6·5
Server-side includes (SSIs), 15·30–15·31, 21·17–21·18
Server Message Block (SMB), 36·4, 36·9
Servers:
buffer overflow attacks, 21·4
dial-up servers and LANs, 25·5–25·7
extranet, 32·17–32·18
Mix servers, 31·2
and mobile code, 17·15–17·16. *See also* Mobile code
proxy servers. *See* Proxy servers
revocation protocols, 37·20–37·21
root servers, 5·25
server spoofing, 28·11–28·12
Web Server security, 21·16–21·19, 27·14
Service-level agreements (SLAs), 47·11, 63·11, 68·28
Service-oriented architecture (SOA), 5·29, 30·24
Service set identifier (SSID), 15·13, 33·7, 33·37
Session border control (SBC), 34·9–34·10
Session hijacking, 36·8, 36·9
Session initiation log records, 53·13
Session Initiation Protocol (SIP), 34·4, 34·9
Session termination log records, 53·13
SET, 21·10
Shamir's *n* out of *m* rule, 37·25
Sharing resources, 24·4–24·5
Shielded twisted pair (STP), 6·10
Shiftwork, 56·31–56·32
Shor's algorithm, 7·41–7·42
Short Message Service (SMS):
and BlackBerrys, 35·16
business threats, 35·14–35·15
overview, 35·3, 35·14, 35·22–35·23
security breach prevention and mitigation, 35·15–35·18
security incident response, 35·18
and VoIP, 31·1
Shoulder surfing, 15·20
Siegel, Martha S., 2·20–2·21
Silver migration, 23·37
Silverlight, 26·16
Simple Certificate Validation Protocol (SCVP), 37·20–37·21
Simple Mail Transfer Protocol (SMTP), 5·28, 17·15, 20·3–20·4, 20·25, 21·7, 21·12
Simple Network Management Protocol (SNMP), 5·23, 5·27, 25·10, 33·41
Simple security property, 24·12
Simultaneous broadcast, 6·6
Single-mode fiber (SMF), 6·11
Site security, 25·3
Site Security Handbook, 44·8–44·9
Site-to-site (S2S) VPNs, 32·7–32·15
Skype, 32·14, 34·2
Slashes, 15·29–15·30
Slowloris attack, 18·12
Small Computer System Interface (SCSI), 57·4
Small to medium business (SMB) appliances, 26·8
Smart Card Security Users Group, 51·17–51·18
Smart cards, 7·31, 7·39, 25·12, 28·2–28·4, 28·15–28·16
Smart Grid Technology Advisory Group, 6·19
Smart phone encryption, 36·15
Smishing, 18·11
Smith, David L., 2·19
Smoke. *See* Fire and smoke
SMURF, 18·9, 25·15
Sniffer Global Analyzer, 33·45
SnifferPro, 25·4
Sniffit, 25·4
Snopes.com, 48·7
Snort, 25·4
Social engineering:
background and history, 19·2–19·4
and computer security incident response, 56·29
consequences of, 19·14–19·15
detection, 19·17–19·19
DNS DoS attacks, 18·10–18·11
Dumpster diving, 19·6–19·7
employee training and awareness, 19·20
examples, 19·15–19·16
frequency of use, 19·4
and malicious code, 16·8
and Kevin Mitnick, 2·5–2·6
low-tech attacks, 19·4, 19·6–19·8
methods, 19·4–19·11
overview, 19·1–19·2
and penetration testing, 46·10–46·11

Social engineering (*Continued*)
pharming. *See* Pharming
phishing. *See* Phishing
physical security, 19·21
prevention of, 19·20–19·21
profile of social engineer, 19·14
psychology, 19·11–19·14
response, 19·19
small business versus large organizations, 19·16–19·17
social psychology, 19·13–19·14
and spam, 48·9
success rate, 19·16
and system penetration, 15·3–15·7
targets of, 19·4–19·5
trends, 19·17
and Trojan horses, 19·2
Social Engineering Defense Architecture (SEDA), 19·19
Social media and social-engineering attacks, 19·10–19·11
Social networking, 48·5, 69·27
Social psychology:
anonymity and aggression, 70·12
anonymity and dishonesty, 70·12
anonymity and prosocial behavior, 70·13–70·14
attribution errors, 50·7–50·10
behavior, explanations of, 50·7
beliefs and attitudes, 50·13–50·16
containment theory, 70·10–70·11
deindividuation theory, 70·9–70·10
general crime theory, 70·11
group behavior, 50·20–50·21
identity in cyberspace, 70·14–70·17
initiative, encouraging, 50·16–50·20
intercultural differences, 50·10–50·11
overview, 50·1–50·2
personality, theories of, 50·4–50·7
rationality, 50·2–50·3
reality, framing, 50·11–50·12
recommendations, 50·22–50·24
reward versus punishment, 50·13
schema, 50·3–50·4
security policies, explaining, 50·12–50·13
self-awareness, 70·12–70·13
in social engineering, 19·13–19·14
social learning theory, 70·11–70·12
and technological generation gap, 50·21–50·22
Social setting leveraging, 19·7–19·8
Society for Worldwide Interbank Funds Transfer (SWIFT), 14·14
Sockets, 5·22–5·23
Software:
commercial off-the-shelf (COTS) software, 14·3, 17·4, 21·13, 21·22, 47·9
component-based software (CBS), 21·13
computer programs, 47·3
content-filtering software, 19·18–19·19
data classification, 67·6–67·8
data integrity, 47·17
downloading, 48·14
dull disk encryption (FDE) software, 40·4
errors, types of, 38·10–38·15, 39·8–39·12
externally supplied, 47·9–4·10
IEEE 802.11 security-auditing software, 33·45–33·46
incompatibilities, 57·26–57·27
and indispensable employees, 45·4–45·6
keys, 42·11–42·12
list server software, 18·7
misconfiguration, 21·17, 21·22
mitigating problems due to denial-of-service attacks, 18·28
new versions of, responsibilities of operations staff, 47·7–47·9
Passfaces, 28·13
password crackers, 15·25–15·26
piracy, 48·31
software inventory management, 40·6
threats to, 24·3
tokens (soft tokens), 28·16–28·17
tracking versions of, 47·7–47·9
usage counters, 42·4–42·5
workplace privacy, 69·23–69·24
Software Assurance Forum for Excellence in Code (SAFECode), 51·14
Software development:
automated testing, 39·16–39·17
bugs, tracking and removal, 39·17
bugs and debugging, 39·5, 39·19–39·20
change management, 39·17–39·18
data corruption, 39·19–39·20
design flaws, 39·19
design phase, 39·6
documentation, 38·13, 39·11, 39·18
errors, types of, 39·8–39·12
hacking, 39·20
implementation flaws, 39·19
implementation phase, 39·6
joint application design (JAD), 39·7
maintenance phase, 39·6
overview, 39·2
rapid application development (RAD), 39·7
regression testing, 39·16, 39·18
requirements analysis, 39·6
security, 66·7–66·8
software development life cycle (SDLC), 39·3–39·8
software quality assurance (SQA), 39·2–39·3, 39·19
test cases, designing, 39·12–39·16
testing, 39·5–39·6, 39·16–39·17
unauthorized changes, 39·20
Websites for secure coding information, 38·15–38·16
Software development life cycle (SDLC):
joint application design (JAD), 39·7
overview, 39·3–39·4
phases, 39·4–39·6

rapid application development (RAD), 39·7
security integration, 39·8
waterfall model, 39·5–39·7
Software Engineering Institute, 44·8, 56·5. *See also* Computer Emergency Response Team Coordination Center (CERT/CC)
Software & Information Industry Association, 42·3
Software patches. *See also* Updates
alternative host architectures, 40·5–40·6
application whitelisting technologies, 40·6–40·7
bundling, 40·3
configuration, 40·4–40·5
enterprise patch management technologies, 40·7–40·10
installation side effects, 40·6
inventory management capabilities, 40·8
metrics and measures, 40·10–40·11
operating system, 47·13–47·14
overview, 40·1–40·2
patch implementation verification, 40·6
patch management capabilities, 40·8
phased deployment, 40·9
prioritizing, 40·4
resource overload, 40·6
software inventory management, 40·6
technology security, 40·9
testing, 40·3–40·4
timing, 40·3–40·4
usability and availability, 40·10
Software total quality management, 38·2–38·4
Solsniff, 25·4
Sony Music, 17·4–17·5
Source libraries, 47·3
Spam:
and anonymity, 70·5–70·6
antispam router, 20·13
CAN-SPAM Act of 2003, 16·11, 20·15, 20·26–20·27
costs of, 20·10–20·12
criminal prosecution, 20·6
defined, 20·6–20·8
email content filtering, 5·28
filters, 20·21–20·24
forged headers, 48·9
global, 2·21
Green Card Lottery, 2·20–2·21
impact of, 20·7–20·8, 20·11–20·13
and IRC bots, 16·11
origin of term, 20·7–20·8
outbound, 20·14
overview, 20·1–20·3
and permissions, 20·16–20·17
preventing, 20·18–20·27
profitability of, 20·8–20·9
recommendations for protecting against, 48·43
scams, 20·10
snowshoe spamming, 20·24
versus SPAM®, 20·7–20·8
SPam over Internet Telephony (SPIT), 34·4–34·5
as unsolicited commercial email (UCE), 20·8
Spanning Tree Protocol (STP), 5·13
Spearphishing, 17·4, 18·11, 33·42, 49·33–49·34
SpectorSoft, 21·9
SPI Dynamics, 21·16
Spim, 19·9
Spoofing:
antispoofing, 26·10
beacon frame spoofing, 33·26–33·27
defined, 8·8
examples of, 8·8
IP address, 5·16, 26·5, 30·33, 31·10–31·11, 33·24–33·25, 36·8, 36·10
MAC address, 33·24–33·25
management frame spoofing, 33·26–33·27
pointer (PTR) record spoofing, 18·10
public key spoofing, 37·4
socket spoofing, 5·22–5·23
spoofed UDP packets, 18·9
Spread spectrum radio transmission, 5·12–5·13
Spybots, 17·13. *See also* Bots
Spyware:
and e-commerce, 21·9
and Internet use, 48·14–48·15, 48·42–48·43
overview, 16·9–16·10
SQL Slammer worm, 16·7–16·8, 53·19
SSH Communications Security, 25·5
SSL/HTTP-based tunnels, 30·35
SSLStrip, 32·12
Standards:
American National Standards Institute (ANSI), 6·23, 37·16
British Standard 7799 (BS 7799), 44·3–44·4, 65·8
Capability Maturity Model (CMM), 51·7, 51·11–51·12
Cloud Security Alliance (CSA), 51·11
Common Criteria (CC), 51·3–51·4, 51·26–51·28
Common Criteria paradigm overviw, 51·28–51·32
and core layers, 5·10
data classification, 67·3–67·6
disaster/emergency preparedness, 22·7
federal government identity cards, 28·17
FedRAMP program, 51·24–51·26
government standard cryptographic validation programs, 51·23–51·24
IEEE. *See* IEEE 802 standards
informal security standards, 44·5–44·9
for information-security governance, 65·8
Internet Engineering Task Force. *See* Internet Engineering Task Force (IETF)
ISO. *See* International Organization for Standardization (ISO)
ISO 9000, 51·7, 51·11
ISO 17799, 38·20
ISO/IEC 15408, 38·20

Standards (*Continued*)
ISO/IEC WD 15443, 38·20
ISO/IEC 17799:2005, 44·3
layered standards architectures, 5·10–5·11
National Electric Code (NEC), 23·8
OSI, 5·11
overview, 44·2–44·3, 51·1–51·2
product development and assessment approaches, 51·9–51·23
Public Key Cryptography Standards (PKCS), 37·16
and risk management, 51·3–51·4
and secure code, 38·20
single network, 5·11–5·12
sources of, 51·5–51·7
SSE-CCM, 38·20
Telecommunications Industry Association/Electronic Industry Alliance (TIA/EIA), 23·8
and trust, 51·3–51·4
types of product-oriented standards, 51·5–51·9
value of, 51·2–51·3
X.509. *See* X.509 certificate format
Star property (*-property), 24·12, 24·14
Star topology, 6·4
Star-wired bus, 6·6–6·7, 6·10
Star-wired ring, 6·6–6·7, 6·10
State governments, data classification policies for, 67·12–67·15
State law:
privacy regulations, 69·27·69·28
and VoIP, 34·3
Stateful inspection, 26·5
Static analysis, 38·16
Static electricity, 4·23
Statistics, and computer crime studies, 10·2–10·9
Steganography, 7·44–7·45, 31·12
Steve Jackson Games, 2·26–2·27
Stop Online Piracy Act (SOPA), 11·41–11·43
Stored Wire and Electronic Communications and Transactional Records Act (SCA), 11·30, 11·32–11·33
Storage area network (SAN), 36·3–36·5, 57·4
Storage media. *See also* Data storage
and data backups. *See* Data backups
and data leakage, 1·18
discarding, 15·19–15·20, 36·15–36·15
nonvolatile media, 36·1
secondary storage, 4·9–4·10
theft, 4·22
volatile devices, 36·2
Storage Networking Industry Association (SNIA), 57·17
Structured Query Language (SQL) injection, 26·4, 26·12, 32·18
Stuxnet worm, 2·20
Subnetwork Access Protocol (SNAP), 6·24
Substitution ciphers, 7·8
Subversion, 17·13, 21·15–21·16
Sun Microsystems, 17·11
Supervisor calls, 4·7
Supervisory control and data acquisition (SCADA), 1·13, 14·9, 53·26–53·27
Surveillance systems, 23·8, 23·11, 23·18–21, 69·14–69·15
Surveys as computer crime research method, 10·6, 10·9–10·10
Switched network services, 5·5
Switches, 5·4–5·5, 5·12, 5·13, 6·25
Symantec, 20·2, 20·11
SYN flooding, 18·11–18·12
Synchronous time, 4·10–4·11
System Access Control List (SACL), 24·17–24·18
System Administration and Network Security (SANS) Institute, 44·10, 51·6, 51·10–51·11
System and asset inventories, 13·8
System boot log record, 53·12
System console activity log record, 53·14
System mode, 24·9
System and network penetration:
exchange of information on, issues with, 15·36
factors, 15·1–15·3
full disclosure, 15·36–15·38
future, 15·39–15·40
nontechnical methods, 15·3–15·7
sources of information on, 15·38–15·39
technical methods, 15·7–15·36
testing, tools, and techniques, 15·20–15·27
through Websites, 15·27–15·31
System shutdown log record, 53·12
System tables, 53·10
Synchronous communications, 15·9
System administrators, password access, 28·5–28·6
System components, monitoring, 53·19

12-Tricks Trojan, 2·13
2600: The Hacker Quarterly, 2·24, 18·5
Tailgating, 19·8
Taxonomy, computer security incident information, 8·4–8·16
TCP (Transmission Control Protocol), 5·18–5·23. *See also* TCP/IP (Transmission Control Protocol/Internet Protocol)
TCP/IP (Transmission Control Protocol/Internet Protocol):
and data communications, 5·9
dynamic routing protocols, 5·27
and e-commerce, 21·10, 30·26
Internet Control Message Protocol (ICMP), 5·24–5·25
and layered standards architectures, 5·10–5·11
and protecting Web applications, 21·6
Telecommunications Industry Association/Electronic Industry Alliance (TIA/EIA), Standard 606, *Administration Standard for the Telecommunications Infrastructure of Commercial Buildings,* 23·8

traffic shaping, 20·25
and UNIX systems, 6·28
TCP port 80, 15·24, 18·23, 30·35
TCP port 443, 30·35, 30·41
TCP three-way handshake, 18·11
Tcpdump, 25·4
Teardrop attacks, 18·10, 25·15, 26·5
Technical knowledge, 4·21
Telecommunications Industry Association/Electronic Industry Alliance (TIA/EIA) Standard 606-B, 23·8
Telecommuting, 1·12
Telephone Consumer Protection Act of 1991, 11·29
Television, piracy, 42·8–42·10
Telnet, 5·28
Temperature, as security hazard, 3·17, 4·11, 4·22
TEMPEST (Transient ElectroMagnetic Pulse Emission Standard), 15·14, 25·5
Temporal key integrity protocol (TKIP), 33·8, 33·9–33·10, 33·16–33·20
Terminals, 4·13–4·14
Terms of use, 11·25–11·26
Terrorism and terrorists. *See also* September 11, 2001, attacks
and hardware backups, 4·18
information infrastructure, protecting against threats, 23·42–23·43
physical threats, 22·4, 22·21–22·22
reporting, 61·5, 61·6–61·7
taxonomy, 8·16
Testing:
automated, 39·16–39·17
race conditions, testing for, 39·14
regression testing, 39·16
software development, 39·5–39·6, 39·12–39·16
state transitions testing, 39·14
test-coverage monitors, 39·14–39·15
test data, 47·15
test libraries, 47·15
Theft, 4·22, 19·7
Theft-of-service attacks, 34·6
Theory of nymity, 70·15
Threat analysis, 3·14–3·17, 22·2, 30·25
Threats:
adaptive threat mitigation, 26·12
classification of damage, 59·3–59·6
environmental, 53·18
health threats, 22·16
healthcare information security, 71·4–71·14
identifying, 22·10–22·11
insider, 13·2–13·8
Internet-based systems, 13·7
list of, 59·2
malicious code threat model, 16·2–16·3
man-made, 22·16–22·18
mitigating, 13·8–13·11
physical threats. *See* Physical threats
reporting, 61·5–61·6
service provider threats, 13·7
system administration threats, 13·8
unified threat management (UTM), 26·4, 26·6
Voice over Internet Protocol (VoIP), 34·4–34·6
wireless local area networks, 33·21–33·24
Threats, assets, vulnerabilities model, 3·2, 3·20–3·22
Thumb drives, 21·9. *See also* Flash drives
Time, synchronous and asynchronous, 4·10–4·11
Time stamps, 47·15
Titan Rain 2·30
TLS. *See* Transport Layer Security (TLS)
Token Bus Working Group, 6·17
Token passing, 6·14–6·15
Token-ring hubs, 6·25
Token ring network, 6·15, 25·5
Token-ring standard, 6·22–6·23
Token Ring Working Group, 6·17
Tokens:
and authentication principles, 28·3–28·4
hardware, 28·16
one-time password generators, 28·17
soft tokens, 28·16–28·17
types of, 28·3–28·4, 28·13
and virtual private networks, 32·6
Toolkits, 8·14, 15·21–15·22
Topology, local area networks, 6·3–6·8
Touch cards, 23·13–23·14
Traceroute (tracert), 5·25
Trade secrets, 11·5–11·8, 11·36–11·37, 45·3
Trademarks, 42·2
Traducement, 9·18
Training:
computer security incident response team, 56·16–56·17
security response plan, 23·45–23·46
Transistors and core memory, 1·7
Transission Control Protocol/Internet Protocol (TCP/IP), 18·2
Transport Layer Security/Secure Sockets Layer (TLS/SSL). *See* Secure Sockets Layer/Transport Layer Security (SSL/TLS)
Transport Layer Security (TLS), 5·26, 30·11–30·12, 32·4, 34·9
Transportation systems and information warfare, 14·11–14·12
Transposition cipher, 7·18–7·20
Trapping, 4·8
Treaties:
intellectual property, 11·35–11·39
Patent Cooperation Treaty (PCT), 11·35
World Intellectual Property Organization Copyright Treaty, 11·15
Trend analysis and reporting, 53·23–53·24
Trespass, 11·25
Trinoo (Trin00), 18·15, 18·19–18·20
Triple A Engine, 70·13
Trojan code, 20·1–20·3, 20·31–20·35

Trojan horses. *See also* Social engineering
antivirus programs
attacks, 20·1–20·3, 20·32–20·34
Back Orifice, 21·4, 21·10
and cyberinvestigation, 55·13–55·14
defenses, 20·34–20·35, 24·14
and malicious Web servers, 16·11
malware, 15·31–15·32
and mobile code, 17·3
overview, 16·8
and passwords, 28·6
and porn sites, 48·35
and screensavers, 20·32
and social engineering attacks, 19·2
system penetration, 15·14–15·15, 15·26
taxonomy, 8·13
Truncated binary exponential backoff, 6·14
Trust:
asymmetric trust, 17·12–17·13
chain of trust, 37·5
derivative trust, 17·12–17·13
domain ranges, trusted and untrusted, 21·7
levels of, 37·10
models, 37·11–37·13
path, 37·15
and PKI interoperability, 37·14–37·15
Pretty Good Privacy (PGP), 7·26, 37·2, 37·13
proofing, 37·10–37·11
and rekeying, 37·22
and signed code, 17·6–17·11. *See also* Mobile code
transitive trust, 17·12–17·13
trusted archival services, 37·25–37·27
trusted paths, 37·10–37·11
trusted time stamp, 37·26–37·27
web of trust, 7·26, 37·5
Trusted archival services, 37·25–37·27
Trusted communication, 30·6–30·7
Trusted time stamp, 37·26–37·27
TRUSTe, 30·20
Trusted Computer Systems Evaluation Criteria (TCSEC) (*Orange Book*), 1·14–1·15, 17·11, 17·15, 25·11, 51·10
Trusted paths, 37·10–37·11
Trusted systems, 24·12–24·14
TSADBOT, 48·14–48·15
Tunneling, 26·10, 26·17, 30·39, 30·41, 31·11, 32·3–32·4
Twisted pair cable, 6·9–6·10, 6·13
Two-factor authentication, 25·6, 28·3, 28·8, 28·17
Typed access control model, 9·6

Unamailer, 2·21–2·22
Unicode, 4·3
Unified Messaging System (UMS), 31·1
Unified Modeling Language (UML), 21·14
Unified threat management (UTM), 26·4, 26·6, 31·9
Uniform resource locator (URL), 31·5, 31·8
Uniform Trade Secrets Act (UTSA), 11·6–11·8
Uninterruptible power supply (UPS), 4·23
United Kingdom, 48·27, 51·6, 69·3
United States Federal Chief Information Officers (CIO) Council:
Best Practices Committee, documentation for policy makers, 44·7–44·8
Universal Automatic Computer (UNIVAC), 1·5
Universal Serial Bus (USB) tokens, 4·9
UNIX, 6·28, 16·10, 17·11, 18·7, 18·20, 21·7, 25·9, 25·12, 25·13–25·14
sendmail program, 16·7, 17·16
Unlimited Operations, 2·28–2·29
Unmanaged hosts, 40·5
Unshielded twisted pair (UTP) wire, 5·3–5·4, 5·12, 6·2, 6·9, 25·5
Updates:
automatic, 16·13, 17·13, 26·22–26·23
importance of, 5·3
Urban myths, 48·7
U.S. Constitution, First Amendment rights, 48·4, 72·12–72·21
U.S. Department of the Navy, 53·7
U.S. Postal Inspection Service, 61·8
U.S. Secret Service, 61·8
USA PATRIOT Act, 69·16, 69·21–69·22, 70·20
USENET, 1·11–1·12, 20·7, 48·4, 48·24, 48·31
User Account Control (UAC), 25·12
User Datagram Protocol (UDP), 5·24, 15·22–15·23, 16·7, 18·9, 18·12, 30·37, 34·4
User identifiers (IDs), 28·2
User interface, 38·13
User virtual machine (UVM), 38·13, 39·10
Utilities, diagnostic utilities, 53·9
UUCP, 30·34

Validation, 4·6, 47·16–47·18
Van Eck phreaking (monitoring devices), 5·12, 15·13–15·14, 25·5
Varonis Systems, 67·7–67·8
Vendors, contracts with, 11·4–11·5
Veracode Vendor Application Security Testing (VAST) program, 51·22
Verification of identity, 29·5–29·7
VeriSign, 7·35, 17·9–17·10, 33·35
Vertical redundancy checks, 4·4
Video:
piracy, 42·8, 48·31
videocassettes, watermarking, 42·11
Video Privacy Protection Act of 1988, 69·21
Vigenère cipher, 7·13–7·14
Violence:
physical threats, 23·42–23·43
prevention and mitigation, 23·7–23·8
threats of, 48·37
video games, 48·30
Virtual firewalls, 26·8–26·9

Virtual local area networks (VLANs), 16·13–16·14, 33·3–33·4, 33·42
Virtual Machine technology, 17·14
Virtual Private Network Consortium, 51·16–51·17
Virtual private networks (VPNs):
  background, 32·1–32·2
  client management, 32·9–32·10
  and e-commerce, 30·11, 30·16
  and encryption, 31·11–31·12, 32·11
  extranets, 32·15–32·19
  information assurance, 32·15–32·16
  infrastructure, 32·6, 32·12–32·13
  interception, 32·11
  and IPv6, 32·14
  and mobile access, 47·16
  network access requirements, 32·7
  overview, 32·2–32·3, 32·20
  processing power, 32·11
  protection, 32·10
  remote access VPNs, 32·4–32·7
  remote-site VPNs, 32·8–32·15
  traffic inspection, 32·11
  and tunneling, 31·11
  and wireless local area network security, 1·19, 25·7–25·8, 26·19
  versus wireless networks, 33·43
Virtual reality, 48·29–48·30
Virtualization and the cloud, 1·13, 53·27–53·29, 57·17–57·18
Viruses:
  boot sector, 16·4
  Brain boot-sector virus, 2·16
  Cascade virus, 2·11, 2·16
  Creeper virus, 2·15
  defined, 16·5
  Elk Cloner virus, 2·15
  financial motivation, 41·2
  email content filtering, 5·28
  hoaxes and Internet myths, 48·7–48·8
  and intrusion detection response, 27·11
  Jerusalem virus (Friday the 13th virus), 2·11, 2·16
  Lehigh virus, 2·16
  MacOS, 25·15
  macro.concept virus, 2·17
  malware, 15·31–15·32, 48·13–48·14
  Melissa virus, 2·18–2·19, 15·37, 18·4, 25·11
  PAPA MS-Excel virus, 18·4
  taxonomy, 8·13, 8·18
  types of, 16·4–16·7
  virus creators, 12·22–12·24
  WildList, 2·18, 2·19, 48·13
Vishing 18·11, 19·10
VisiCalc spreadsheet program, 1·8, 1·9–1·10
Visitors, 47·6–47·7
Visitors, controlling, 47·6–47·7
Visual authentication, 23·14
Visual Basic for Applications (VBA), 16·4
Visual BASIC (VB), 47·3
Voice over Internet Protocol (VoIP):
  application layer gateways and firewalls, 34·7
  and application security, 5·28
  audio stream protocols, 34·4
  device authentication, 34·8
  encryption, 34·9–34·10
  logical separation of voice and data, 34·7
  network monitoring tools, 34·7–34·8
  overview, 34·1–34·2
  quality of service (QOS), 34·7
  real-time antivirus scanning, 34·7
  signaling protocols, 34·4
  threats, 34·4–34·6
  user authentication, 34·8
  and user datagram protocol, 5·23, 5·24
  vishing, 18·11, 19·10
Voice phishing, 19·10
Voice-print scanning, 23·15
Vulnerability:
  allowed path vulnerabilities, 21·7
  analysis tools, 15·21–15·22
  assessment, 26·24, 46·2–46·7
  class analysis, 15·29
  credentialed monitoring, 46·6
  healthcare information security, 71·4–71·14
  management, 46·1–46·3
  and malware, 41·2
  noncredentialed monitoring, 46·6
  scanning, 46·4, 55·14
  and security incident common language, 8·11–8·15
  and segmented secrets, 23·5–23·6
  and software patches, 40·1. *See also* Software patches
Vulnerability Discovery ad Remediation Open Source Hardening Project, 51·18
Vulnerability scanning, 38·19

WANK worm, 30·39
War (demon) dialing, 15·16, 21·12, 46·4
War driving, 15·13, 15·20, 46·4
Water damage, 4·12
Watermarking, 42·11–42·12
Weapons of mass destruction (WMDs), 22·21–22·22, 23·42–23·43
Weather as security threat, 22·12
Web 2.0, 35·18
Web application system security, 21·5–21·8, 26·4
Web beacons (bugs), 69·25
Web browsers, 1·12, 15·18, 16·13, 17·13
Web crawlers, 11·23, 11·28
Web monitoring:
  anonymity and privacy concerns, 31·12
  block lists, 31·9
  and caching services, 31·13
  and encryption, 31·11–31·12
  filtering methods, 31·4–31·8

Web monitoring (*Continued*)
firewalls, 31·9–31·10
implementation, 31·8–31·9
and IP spoofing, 31·10–31·11
and language translation sites, 31·12–31·13
national security surveillance, 31·4
need for, 31·1–31·2
overview, 31·14
parental tools, 31·10, 31·13
and pornography, 48·36
and privacy issues, 69·25–69·26
proxy servers, 31·2, 31·9, 31·13
reasons for, 31·2–31·4
terminology, 31·2
trends, 31·13
tunneling, 31·11
vulnerabilities, 31·10–31·13
Web proxy, 16·13
Web servers, 16·11, 21·16–21·19
Web of trust, 7·26, 37·5
Websites:
customer monitoring, 30·42–30·43
maintenance and updating, 30·32
personal, 48·5
privacy policies, 30·20
protection of and e-commerce, 30·18–30·22
secure coding resources, 38·15–38·16
security, 30·32–30·33
Snopes.com, 48·7
system and network penetration through Websites, 15·27–15·31
term paper analysis, 48·32
Web page defacement, 15·28, 15·30
WildList, 2·18, 2·19, 41·5, 48·13
WebInspect, 21·16, 39·14
WebSockets, 17·3
Wellenreiter, 33·46
Whaling, 17·4
White box testing, 38·17
Wide area networks (WANs), 1·12, 5·5–5·6, 6·2, 32·7–32·8
Wi-Fi, 4·16, 15·12–15·13, 35·18
Wi-Fi Protected Access (WPA), 5·3
Wi-Fi Protected Setup (WPS), 33·32–33·33
WikiLeaks, 70·6
WildList, 2·18, 2·19, 41·5, 48·13
Window field, 5·21
Windows. *See* Microsoft Windows
Windows wireless client, 33·46
WinDump, 25·4
Wired Equivalent Privacy (WEP), 25·7–25·8, 33·4, 33·8
Wired facilities, 4·14–4·16
Wireless access point (WAP), 5·3
Wireless Coexistence Working Group, 6·19
Wireless communications, 4·16
Wireless devices, 6·2, 6·11–6·13, 21·8
Wireless intrusion-detection systems (WIDS), 33·37–33·40
Wireless LAN (WLAN) Working Group, 6·18
Wireless local area network (WLAN):
attacks, 33·24–33·34
authentication and access control, 33·7–33·8
benefits of, 33·3–33·4
corporate use of, 33·2
data confidentiality, 33·8
fast secure roaming, 33·20
IEEE 802.11 Robust Security Network, 33·8, 33·9–33·20
IEEE 802.11 security fundamentals, 33·6–33·9
IEEE 802.11 standards, 33·5–33·6. *See also* IEEE 802 standards
intrusion detection and prevention systems (WIDPS), 27·14
key management, 33·8–33·9
management, 33·4
mitigating controls, 33·34–33·40
network penetration techniques, 15·12–15·13
overview, 32·2–32·3
secure enterprise design, 33·40–33·44
security auditing tools, 33·44–33·46
security issues, 25·7–25·9
threats, 33·21–33·24
Wireless networks, 5·3, 5·12, 15·12–15·13, 15·20
Wireless Personal Area Network (WPAN) Working Group, 6·18
Wireless phones, 15·12
Wireless Regional Area Networks (WRAN) Working Group, 6·19
Wireshark, 33·46
Wiretap Act, 11·30–11·32
Wiretapping, 5·4, 5·12, 15·9, 22·18–22·19, 23·41–23·42
Workplace violence, 22·3, 22·21–22·22
World Intellectual Property Organization Copyright Treaty, 11·15
World Trade Organization (WTO) Agreement, 11·35
World Wide Name (WWN) service, 36·8
World Wide Web, 1·12–1·13
Worms:
Bagel worm, 16·8
Christmas Tree worm, 2·16, 18·3, 18·7
Code Red Worm, 18·23–18·24
and cyberinvestigation, 55·13
first worm, 1·9
Internet Worm, 18·7
and intrusion detection response, 27·11
"Love Bug" worm, 2·20
Love Letter (ILOVEYOU) worm, 18·4, 25·11
MacOS, 25·15
and malicious code, 16·7–16·8
malware, 15·31–15·32
Morris Worm, 2·16–2·17, 16·7, 18·3, 18·5, 30·39, 56·4–56·5

NIMDA worm, 16·7, 18·24–18·25
PrettyPark worm, 16·10
SQL Slammer, 16·7–16·8
Stuxnet worm, 2·20
taxonomy, 8·18
WANK worm, 30·39
Warhol Worms, 16·7
Wozniak, Steve, 1·8, 2·8
WPA/WPA2 pre-shared key, 33·31–33·32

X.25 carriers, 15·10
X.509 certificate format, 7·31–7·35, 17·6, 30·39, 33·35, 37·5–37·6, 37·8, 37·16–37·20
Xerox:
antipiracy programs, 42·5
and digital rights management, 42·14
XOR (Exclusive-Or), 7·15–7·16, 53·16

Y2K problem, 2·11–2·12
Yahoo!, 18·15, 48·22

zBubbbles, 48·14
Zero-day attacks, 15·25
Zero latency, 30·24
ZigBee, 53·17–53·18
ZoneAlarm, 48·15, 48·43
Zoned architecture, 26·18